arXiv:1602.03124v4 [cs.CC] 27 Nov 2016

Even Delta-Matroids and the Complexity of Planar Boolean CSPs*

Alexandr Kazda† Vladimir Kolmogorov‡ Michal Rolínek§

Abstract

The main result of this paper is a generalization of the classical blossom algorithm for finding perfect matchings. Our algorithm can efficiently solve Boolean CSPs where each variable appears in exactly two constraints (we call it edge CSP) and all constraints are *even* Δ*-matroid* relations (represented by lists of tuples). As a consequence of this, we settle the complexity classification of planar Boolean CSPs started by Dvořák and Kupec.

Knowing that edge CSP is tractable for even Δ-matroid constraints allows us to extend the tractability result to a larger class of Δ-matroids that includes many classes that were known to be tractable before, namely *co-independent*, *compact*, *local* and *binary*.

1 Introduction

The constraint satisfaction problem (CSP) has been a classical topic in computer science for decades. Aside from its indisputable practical importance, it has also heavily influenced theoretical research. The uncovered connections between CSP and areas such as graph theory, logic, group theory, universal algebra, or submodular functions provide some striking examples of the interplay between CSP theory and practice.

We can exhibit such connections especially if we narrow our interest down to *fixed-template CSPs*, that is, to sets of constraint satisfaction instances in which the constraints come from a fixed set of relations Γ. For any fixed Γ the set of instances CSP(Γ) forms a decision problem; the question if CSP(Γ) is always either polynomial-time solvable or NP-complete (in other words it avoids intermediate complexities assuming $P \neq NP$) is known as the CSP dichotomy conjecture [12] and is one of more notorious open problems in theoretical computer science.

In this work we address two special structural restrictions for CSPs with Boolean variables. One is limiting to at most two constraints per variable and the other requires the constraint network to have a planar representation. The first type, introduced by Feder [10], has very natural interpretation as CSPs in which edges play the role of variables and nodes the role of constraints, which is why we choose to refer to it as *edge CSP*. It was Feder who showed the following hardness result: Assume the constraint language Γ contains both unary constant relations. Then unless all relations in Γ are Δ-matroids, the edge CSP with constraint language Γ has the same complexity as the unrestricted CSP with constraint language Γ. Since then, there has also been progress on the algorithmic side. Several tractable (in the sense of being polynomial time solvable) classes of Δ-matroids were identified [10, 11, 15, 7, 13, 8]. A recurring theme is the connection between Δ-matroids and matching problems.

*This work was supported by European Research Council under the European Unions Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement no 616160.

†IST Austria, Am Campus 1, 3400, Klosterneuburg, Austria, alex.kazda@gmail.com

‡IST Austria, Am Campus 1, 3400, Klosterneuburg, Austria, vnk@ist.ac.at

§IST Austria, Am Campus 1, 3400, Klosterneuburg, Austria, michalrolinek@gmail.com

Recently, a setting for planar CSPs was formalized by Dvořák and Kupec [8]. In their work, they provide certain hardness results together with a reduction of the remaining cases to Boolean edge CSP. Dvořák and Kupec's results imply that completing the complexity classification of Boolean planar CSPs is equivalent to establishing the complexity of (planar) Boolean edge CSP where all the constraints are *even* Δ-*matroids*. In their paper, Dvořák and Kupec provided a tractable subclass of even Δ-matroids along with computer-aided evidence that the subclass (matching realizable even Δ-matroids) covers all even Δ-matroids of arity at most 5. However, it turns out that there exist even Δ-matroids of arity 6 that are not matching realizable; we provide an example of such a Δ-matroid in Appendix B.

The main result of our paper is a generalization of the classical Edmonds' blossom-shrinking algorithm for matchings [9] that we use to efficiently solve edge CSPs with even Δ-matroid constraints. This settles the complexity classification of planar CSP. Moreover, we give an extension of the algorithm to cover a wider class of Δ-matroids. This extension subsumes (to our best knowledge) all previously known tractable classes.

The paper is organized as follows. In the introductory Sections 2, 3, and 4 we formalize the frameworks, discuss how dichotomy for Boolean planar CSP follows from our main theorem and sharpen our intuition by highlighting similarities between edge-CSPs and perfect matching problems, respectively. The algorithm is described in Section 5 and the proofs required for showing its correctness are in Appendix A. The extension of the algorithm is discussed in Appendices C and D.

2 Preliminaries

Definition 1. A Boolean CSP *instance* I is a pair $(V, \mathcal{C})$ where V is the set of *variables* and $\mathcal{C}$ the set of *constraints* of I. A k-ary constraint $C \in \mathcal{C}$ is a pair (σ, R_C) where $\sigma \subseteq V$ is a set of size k (called the *scope* of C) and $R_C \subseteq \{0,1\}^\sigma$ is a relation on $\{0,1\}$. A solution to I is a mapping $\hat{f} : V \to \{0,1\}$ such that for every constraint $C = (\sigma, R_C) \in \mathcal{C}$, $\hat{f}$ restricted to σ lies in R_C.

If all constraint relations of I come from a set of relations Γ (called the *constraint language*), we say that I is a Γ-instance. For Γ fixed, we will denote the problem of deciding if a Γ-instance given on input has a solution by CSP(Γ).

Note that the above definition is not fully general in the sense that it *does not allow one variable to occur multiple times in a constraint*; we have chosen to define Boolean CSP in this way to make our notation a bit simpler. This can be done without loss of generality as long as Γ contains the equality constraint (i.e. $\{(0,0),(1,1)\}$): If a variable, say v, occurs in a constraint multiple times, we can add extra copies of v to our instance and join them together by the equality constraint to obtain a slightly larger instance that satisfies our definition.

For brevity of notation, we will often not distinguish a constraint $C \in \mathcal{C}$ from its constraint relation R_C; the exact meaning of C will always be clear from the context. Even though in principle different constraints can have the same constraint relation, our notation would get cumbersome if we wrote R_C everywhere.

The main point of interest is classifying the computational complexity of CSP(Γ). Constraints of an instance are specified by lists of tuples in the corresponding relations and thus those lists are considered to be part of the input. We will say that Γ contains the unary constant relations if $\{(0)\},\{(1)\} \in \Gamma$ (these relations allow us to fix the value of a certain variable to 0 or 1).

For Boolean CSPs (where variables are assigned Boolean values), the complexity classification of CSP(Γ) due to Schaefer has been known for a long time [17]. There also exists a classification

for the three-element domain [3] and a substantial progress is being made on the general case (see [4, 1, 5]). However, in this work we concentrate on Boolean domains only.

Our main focus is on restricted forms of the CSP. In particular, we are interested in structural restriction, i.e. in restriction on the constraint network. Once one starts to limit the shape of instances, the Boolean case becomes complicated again. (As a side note, we expect similar problems for larger domains to be very hard to classify. For example, Dvořák and Kupec note that one can encode coloring a planar graphs by four colors as a class of planar CSPs that always have a solution for a highly nontrivial reason, namely the four color theorem.)

A natural structural restriction would be to limit the number of constraints in whose scope a variable can lie. When $k \geq 3$ and Γ contains all unary constants, then $\mathrm{CSP}(\Gamma)$ with each variable in at most k constraints is polynomial time equivalent to unrestricted $\mathrm{CSP}(\Gamma)$, see [7, Theorem 2.3]. This leaves instances with at most two occurrences per variable in the spotlight. To make our arguments clearer, we will assume that each variable occurs exactly in two constraints (following [10], we can reduce decision CSP instances with at most two appearances of each variable to instances with exactly two appearances by taking two copies of the instance and identifying both copies of v whenever v is a variable that originally appeared in only a single constraint).

Definition 2 (Edge CSP)**.** Let Γ be a constraint language. Then the problem $\mathrm{CSP}_{\mathrm{EDGE}}(\Gamma)$ is the restriction of $\mathrm{CSP}(\Gamma)$ to those instances in which every variable is present in exactly two constraints.

Perhaps a more natural way to look at an instance I of an edge CSP is to consider a graph whose edges correspond to variables of I and nodes to constraints of I. Constraints (nodes) are incident with variables (edges) they interact with. In this (multi)graph, we are looking for a satisfying Boolean edge labeling. Viewed like this, edge CSP becomes a counterpart to the usual CSP where variables are typically identified with nodes and constraints with (hyper)edges. The idea of "switching" the role of (hyper)edges and vertices already appeared in the counting CSP community under the name Holant problems [6].

This type of CSP is sometimes called "binary CSP" in the literature [8]. However, this term is very commonly used for CSPs whose all constraints have arity at most two [18]. In order to resolve this confusion (and for the reasons described in the previous paragraph), we propose the term "edge CSP".

As we said above, we we will only consider Boolean edge CSP, often omitting the word "Boolean" for space reasons. The following Boolean-specific definitions will be useful to us:

Definition 3. Let $f\colon V \to \{0,1\}$ and $v \in V$. We will denote by $f \oplus v$ the mapping $V \to \{0,1\}$ that agrees with f on $V \setminus \{v\}$ and has value $1 - f(v)$ on v. For a set $S = \{s_1, \ldots, s_k\} \subseteq V$ we let $f \oplus S = f \oplus s_1 \oplus \cdots \oplus s_k$. Also for $f, g\colon V \to \{0,1\}$ let $f \,\Delta\, g \subseteq V$ be the set of variables v for which $f(v) \neq g(v)$.

Definition 4. Let V be a set. A nonempty subset M of $\{0,1\}^V$ is called a Δ-*matroid* if whenever $f, g \in M$ and $v \in f \,\Delta\, g$, then there exists $u \in f \,\Delta\, g$ such that $f \oplus \{u, v\} \in M$. If moreover, the parity of the number of ones over all tuples of M is constant, we have an *even* Δ-*matroid* (note that in that case we never have $u = v$ so $f \oplus \{u, v\}$ reduces to $f \oplus u \oplus v$).

The strongest hardness result on Boolean edge CSP is from Feder.

Theorem 5 ([10])**.** *If Γ is a constraint language containing unary constant relations such that $\mathrm{CSP}(\Gamma)$ is NP-Hard and there is $R \in \Gamma$ which is not a Δ-matroid, then $\mathrm{CSP}_{\mathrm{EDGE}}(\Gamma)$ is NP-Hard.*

Tractability was shown for special classes of Δ-matroids, namely binary [13, 7], co-independent [10], compact [15], and local [7] (see the definitions in the respective papers). All the proposed

algorithms are based on variants of searching for augmenting paths. In this work we propose a more general algorithm that involves both augmentations and contractions. In particular, we prove the following.

Theorem 6. *If Γ contains only even Δ-matroid relations, then $\mathrm{CSP}_{\mathrm{EDGE}}(\Gamma)$ can be solved in polynomial time.*

Our algorithm will in fact be able to solve even a certain optimization version of the edge CSP (corresponding to finding a maximum matching). This is discussed in detail in Section 5.

In Appendix C we show that if a class of Δ-matroids is *effectively coverable*, then it defines a tractable CSP. The whole construction is similar to, but more general than, $\mathcal{C}$-zebra Δ-matroids introduced in [11]. We note here also that the class of coverable Δ-matroids is natural in the sense of being closed under direct products and identifying variables (in other words, gadget constructions).

Definition 7. Let M be a Δ-matroid. We say that $\alpha, \beta \in M$ are *even-neighbors* if there exist distinct variables $u, v \in V$ such that $\beta = \alpha \oplus u \oplus v$ and $\alpha \oplus u \notin M$. We say we can reach $\gamma \in M$ from $\alpha \in M$ if there is a chain $\alpha = \beta_0, \beta_1, \ldots, \beta_n = \gamma$ where each pair β_i, β_{i+1} are even-neighbors.

Definition 8. We say that M is *coverable* if for every $\alpha \in M$ there exists M_α such that:

(a) M_α is an even Δ-matroid (over the same ground set as M),

(b) M_α contains all $\beta \in M$ that can be reached from α (including α itself),

(c) whenever $\gamma \in M$ can be reached from α and $\gamma \oplus u \oplus v \in M_\alpha \setminus M$, then $\gamma \oplus u, \gamma \oplus v \in M$.

In our algorithm, we will need to have access to the sets M_α, so we need to assume that all our Δ-matroids, in addition to being coverable, come from a class of Δ-matroids where the sets M_α can be determined quickly. This is what *effectively coverable* means (for a formal definition see Definition 43).

The following theorem is a strengthening of a result from [11]:

Theorem 9. *Given an edge CSP instance I with effectively coverable Δ-matroid constraints, an optimal edge labeling f of I can be found in time polynomial in $|I|$. In particular, $\mathrm{CSP}_{\mathrm{EDGE}}(\Gamma)$ can be solved in polynomial time.*

As we show in Appendix D, effectively coverable Δ-matroid classes include numerous known tractable classes of Δ-matroids: $\mathcal{C}$-zebra Δ-matroids [11] for any $\mathcal{C}$ subclass of even Δ-matroids (where we assume, just like in [11], that we are given the zebra representations on input) as well as co-independent [10], compact [15], local [7], and binary [13, 7] Δ-matroids. To our best knowledge these are all the known tractable classes and according to [7] the classes other than $\mathcal{C}$-zebras are pairwise incomparable.

3 Implications

In this section we explain how our result implies full complexity classification of planar Boolean CSPs.

Definition 10. Let Γ be a constraint language. Then $\mathrm{CSP}_{\mathrm{PLANAR}}(\Gamma)$ is the restriction of $\mathrm{CSP}(\Gamma)$ to the set of instances for which there exists a planar graph $G(V, E)$ such that $v_1, \ldots, v_k$ is a face of G (with nodes listed in counter-clockwise order) if and only if there is a unique constraint imposed on the tuple of variables $(v_1, \ldots, v_k)$.

It is also noted in [8] that checking whether an instance has a planar representation can be done efficiently (see e.g.. [14]) and hence it does not matter if we are given a planar drawing of G as a part of the input or not. The planar restriction does lead to new tractable cases, for example planar NAE-3-SAT (Not-All-Equal 3-Satisfiability) [16].

Definition 11. A relation R is called *self-complementary* if for all $T \in \{0,1\}^n$ we have $T \in R$ if and only if $T \oplus (1,\ldots,1) \in R$.

Definition 12. For a tuple of Boolean variables $T = (x_1,\ldots,x_n)$, let $dT = (x_1 \oplus x_2,\ldots,x_n \oplus x_1)$. For R relation and Γ set of relations, let $dR = \{dT\colon T \in R\}$ and $d\Gamma = \{dR\colon R \in \Gamma\}$.

Since self-complementary relations don't change when we flip all their coordinates, we can describe a self-complementary relation by looking at the differences of neighboring coordinates; this is exactly the meaning of dR. Note that these differences are realized over edges of the given planar graph.

Knowing this, it is not so difficult to imagine that via switching to the planar dual of G, one can reduce a planar CSP instance to some sort of edge CSP instance. This is in fact part of the following theorem from [8]:

Theorem 13. *Let Γ be such that $\mathrm{CSP}(\Gamma)$ is NP-Hard. Then:*

(a) If there is $R \in \Gamma$ that is not self-complementary, then $\mathrm{CSP}_{\mathrm{PLANAR}}(\Gamma)$ is NP-Hard.

(b) If every $R \in \Gamma$ is self-complementary and there exists $R \in \Gamma$ such that dR is not even Δ-matroid, then $\mathrm{CSP}_{\mathrm{PLANAR}}(\Gamma)$ is NP-Hard.

(c) If every $R \in \Gamma$ is self-complementary and dR is an even Δ-matroid, then $\mathrm{CSP}_{\mathrm{PLANAR}}(\Gamma)$ is polynomial-time reducible to

$$\mathrm{CSP}_{\mathrm{EDGE}}(d\Gamma \cup \{EVEN_1, EVEN_2, EVEN_3\})$$

where $EVEN_i = \{(x_1,\ldots x_i)\colon x_1 \oplus \cdots \oplus x_i = 0\}$.

Using Theorem 6, we can finish this classification:

Theorem 14 (Dichotomy for $\mathrm{CSP}_{\mathrm{PLANAR}}$)**.** *Let Γ be a constraint language. Then $\mathrm{CSP}_{\mathrm{PLANAR}}(\Gamma)$ is solvable in polynomial time if either*

(a) $\mathrm{CSP}(\Gamma)$ is solvable in polynomial time or;

(b) Γ contains only self-complementary relations R such that dR is an even Δ-matroid.

Otherwise, $\mathrm{CSP}_{\mathrm{PLANAR}}(\Gamma)$ is NP-Hard.

Proof. By Theorem 13 the only unresolved case reduces to solving

$$\mathrm{CSP}_{\mathrm{EDGE}}(d\Gamma \cup \{EVEN_1, EVEN_2, EVEN_3\}).$$

Since the relations $EVEN_i$ are even Δ-matroids for every i, this is polynomial-time solvable thanks to Theorem 6. □

4 Even Δ-matroids and matchings

In this section we highlight the similarities and dissimilarities between even Δ-matroid CSPs and matching problems. These similarities will guide us on our way through the rest of the paper.

Example 15. For $n \in \mathbb{N}$ consider the "perfect matching" relation $M_n \subseteq \{0,1\}^n$ containing precisely the tuples in which exactly one coordinate is set to one and all others to zero. Note that M_n is an even Δ-matroid for all n. Then the instance I of $\mathrm{CSP}_{\mathrm{EDGE}}(\{M_n \colon n \in \mathbb{N}\})$ (represented in Figure 1) is equivalent to deciding whether the graph of the instance has a perfect matching (every node is adjacent to precisely one edge with label 1).

One may also construct an equivalent instance I' by "merging" some parts of the graph (in Figure 1 those are X and Y) to single constraint nodes. The constraint relations imposed on the "supernodes" record sets of outgoing edges which can be extended to a perfect matching on the subgraph induced by the "supernode". For example, in the instance I' the constraints imposed on X and Y would be (with variables ordered as in Figure 1):

$$X = \{10000, 01000, 00100, 00010, 11001, 10101, 10011\},$$
$$Y = \{001, 010, 100, 111\}.$$

It is easy to check that both X and Y are even Δ-matroids.

One takeaway from this example is that any algorithm that solves edge CSP for the even Δ-matroid case has to work for perfect matchings in graphs as well. Another is the construction of even Δ-matroids X and Y which can be generalized as follows.

Definition 16 (Matching realizable relations)**.** Let G be a graph and let $v_1, \ldots, v_a \in V(G)$ be distinct nodes of G. For an a-tuple $T = (x_1, \ldots, x_a) \in \{0,1\}^a$, we denote by G_T the graph obtained from G by deleting all nodes v_i such that $x_i = 1$. Then we can define

$$M(G, v_1, \ldots, v_a) = \{T \in \{0,1\}^a \colon \\ G_T \text{ has a perfect matching}\}.$$

We say that a relation $R \in \{0,1\}^a$ is *matching realizable* if $R = M(G, v_1, \ldots, v_a)$ for some graph G and nodes $v_1, \ldots, v_a \in V(G)$.

Every matching realizable relation is an even Δ-matroid [8]. Also, it should be clear from the definition and the preceding example that $\mathrm{CSP}_{\mathrm{EDGE}}(\Gamma)$ is tractable if Γ contains only matching realizable relations (assuming we know the graph G and the nodes $v_1, \ldots, v_a$ for each relation): One can simply replace each constraint node with the corresponding graph and then test for existence of perfect matching.

The authors of [8] also verify that every even Δ-matroid of arity at most 5 is matching realizable. However, as we prove in Appendix B, this is not true for higher arities.

Proposition 17. *There exists an even Δ-matroid of arity 6 which is not matching realizable.*

Proposition 17 shows that we cannot hope to simply replace the constraint nodes by graphs and run the Edmonds' algorithm. The Δ-matroid constraints can exhibit new and more complicated behavior than just matchings in graphs, as we shall soon see. In fact, there is a known exponential lower bound for solving edge CSP with matroid constraints (matroids being special cases of even Δ-matroids) given by oracles (i.e. not lists of tuples) [2], which rules out any polynomial time algorithm that would work in the oracle model. In particular, we are convinced that our method of contracting blossoms can not be significantly simplified while still staying polynomial time computable.

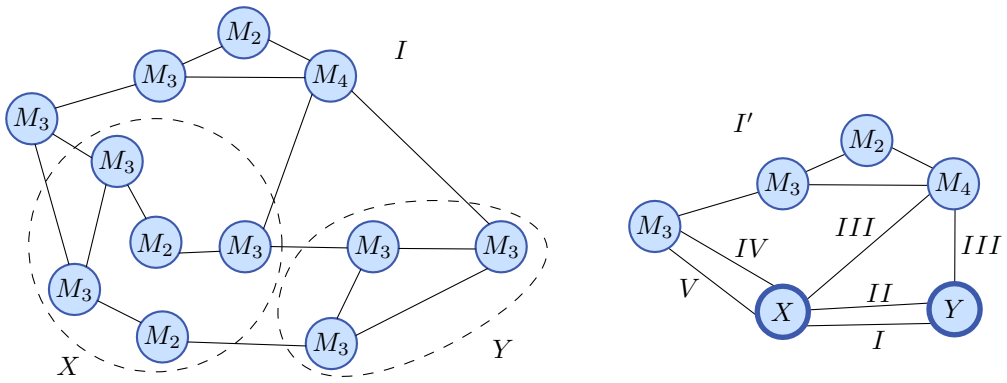

Figure 1: On the left we see an instance I that is equivalent to testing for perfect matching of the given graph. On the right is an equivalent instance I' with contracted "supernodes" X and Y.

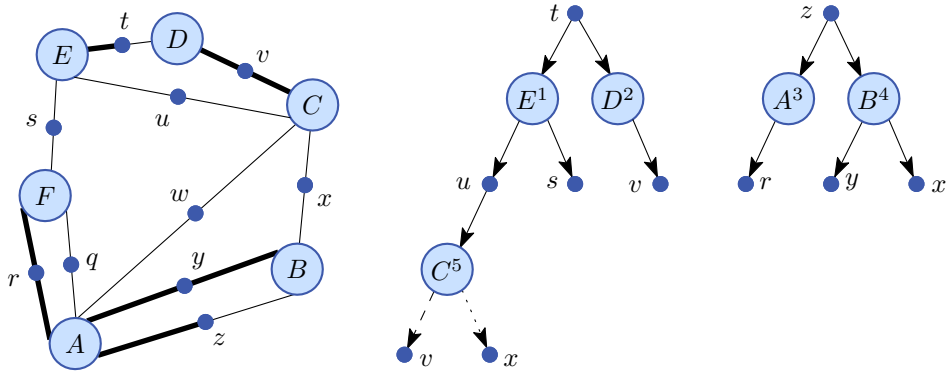

Figure 2: A possible run of Algorithm 1 on the instance I' from Example 15 (with renamed constraint nodes) where the edge labeling f is marked by thick (1) and thin (0) half-edges. We see that the algorithm finds a blossom when it hits the variable v the second time in the same tree. However, had we first processed the transition Cx (which we could have done), we would have found an augmenting path $p = \text{walk}(C^5)\,\text{walk}(x)^{-1}$ (where $\text{walk}(x)^{-1}$ ends in z).

5 Algorithm

5.1 Setup

We can draw edge CSP instances as constraint graphs: The *constraint graph* $G_I = (V \cup \mathcal{C}, \mathcal{E})$ of I is a bipartite graph with partitions V and $\mathcal{C}$. There is an edge $\{v, C\} \in \mathcal{E}$ if and only if v belongs to the scope of C. Throughout the rest of the paper we use lower-case letters for variable nodes in V $(u, v, x, y, \ldots)$ and upper-case letters for constraint nodes in $\mathcal{C}$ $(A, B, C, \ldots)$. Since we are dealing with edge CSP, the degree of each node $v \in V$ in G_I is exactly two and since we don't allow a variable to appear in a constraint twice, G_I has no multiple edges. For such instances I we introduce the following terminology and notation.

Definition 18. An *edge labeling* of I is a mapping $f : \mathcal{E} \to \{0, 1\}$. For a constraint $C \in \mathcal{C}$ with the scope σ we will denote by $f(C)$ the tuple in $\{0, 1\}^\sigma$ such that $f(C)(v) = f(\{v, C\})$ for all $v \in \sigma$. Edge labeling f will be called *valid* if $f(C) \in C$ for all $C \in \mathcal{C}$.

Variable $v \in V$ is called *consistent* in f if $f(\{v, A\}) = f(\{v, B\})$ for the two distinct edges $\{v, A\}, \{v, B\} \in \mathcal{E}$ of G_I. Otherwise, v is *inconsistent* in f.

A valid edge labeling f is *optimal* if its number of inconsistent variables is minimal among all valid edge labelings of I. Otherwise f is called *non-optimal*.

Note that I has a solution if and only if an optimal edge labeling f of I has no inconsistent variables.

The main theorem we prove is the following strengthening of Theorem 6.

Theorem 19. *Given an edge CSP instance I with even Δ-matroid constraints, an optimal edge labeling f of I can be found in time polynomial in $|I|$.*

Walks and blossoms When studying matchings in a graph, paths and augmenting paths are important. We will use analogous objects, called f-walks and augmenting f-walks, respectively.

Definition 20. A *walk* q of length k in the instance I is a sequence $q_0 C_1 q_1 C_2 \dots C_k q_k$ where the variables q_{i-1}, q_i lie in the scope of the constraint C_i, and each edge $\{v, C\} \in \mathcal{E}$ is traversed at most once: vC and Cv occur in q at most once, and they do not occur simultaneously.

Note that q can be viewed as a walk in the graph G_I that starts and ends at nodes in V. Since each node $v \in V$ has degree two in G_I, the definitions imply that v can be visited by q at most once, with a single exception: we may have $q_0 = q_k = v$, with $q = vC \dots Dv$ where $C \neq D$. We allow walks of length 0 (i.e. single vertex walks) for formal reasons.

A *subwalk* of q, denoted by $q_{[i,j]}$, is the walk $q_i C_{i+1} \dots C_j q_j$ (again, we need to start and end in a variable). The inverse walk to q, denoted by q^{-1}, is the sequence $q_k C_k \dots q_1 C_1 q_0$. Given two walks p and q such that the last node of p is the first node of q, we define their concatenation pq in the natural way. If $p = \alpha_1 \dots \alpha_k$ and $q = \beta_1 \dots \beta_\ell$ are sequences of nodes of a graph where α_k and β_1 are different but adjacent, we will denote the sequence $\alpha_1 \dots \alpha_k \beta_1 \dots \beta_\ell$ also by pq (or sometimes as p, q).

If f is an edge labeling of I and q a walk in I, we denote by $f \oplus q$ the mapping that takes f and flips the values on all variable-constraint edges encountered in q, i.e.

$$(f \oplus q)(\{v, C\}) = \begin{cases} 1 - f(\{v, C\}) & \text{if } q \text{ contains } vC \text{ or } Cv \\ f(\{v, C\}) & \text{otherwise} \end{cases} \tag{1}$$

Definition 21. Let f be a valid edge labeling of instance I. A walk $q = q_0 C_1 q_1 C_2 \dots C_k q_k$ with $q_0 \neq q_k$ will be called an *f-walk* if

(a) variables $q_1, \dots, q_{k-1}$ are consistent in f, and

(b) $f \oplus q_{[0,i]}$ is a valid edge labeling for any $i \in [1, k]$.

If in addition variables q_0 and q_k are inconsistent in f then q will be called an *augmenting f-walk*.

Later we will show that a valid edge labeling f is non-optimal if and only if there exists an augmenting f-walk. Note that one direction is straightforward: If p is an augmenting f-walk, then $f \oplus p$ is valid and has 2 fewer inconsistent variables than f.

Another structure used by the Edmonds' algorithm for matchings is a *blossom*. The precise definition of a blossom in our setting (Definition 36) is a bit technical. Informally, an f-blossom is a walk $b = b_0 C_1 b_1 C_2 \dots C_k b_k$ with $b_0 = b_k$ such that:

(a) variable $b_0 = b_k$ is inconsistent in f while variables $b_1, \dots, b_{k-1}$ are consistent, and

(b) $f \oplus b_{[i,j]}$ is a valid edge labeling for any non-empty proper subinterval $[i,j] \subsetneq [0,k]$,

(c) there are no bad shortcuts inside b (we will make this precise later).

5.2 Algorithm description

We are given an instance I of edge CSP with even Δ-matroid constraints together with an edge labeling f and we want to either show that f is optimal or improve it. Our algorithm will explore the graph $(V \cup \mathcal{C}, \mathcal{E})$ building a directed forest T. Each variable node $v \in V$ will be added to T at most once. Constraint nodes $C \in \mathcal{C}$, however, can be added to T multiple times. To tell the copies of C apart (and to keep track of the order in which we built T), we will mark each C with a timestamp $t \in \mathbb{N}$; the resulting node of T will be denoted as $C^t \in \mathcal{C} \times \mathbb{N}$. Thus, the forest will have the form $T = (V(T) \cup \mathcal{C}(T), E(T))$ where $V(T) \subseteq V$ and $\mathcal{C}(T) \subseteq \mathcal{C} \times \mathbb{N}$.

The roots of the forest T constructed by the algorithm will be the inconsistent nodes of the instance (for current f); all non-root nodes in $V(T)$ will be consistent. The edges of T will be oriented towards the leaves. Thus, each non-root node $\alpha \in V(T) \cup \mathcal{C}(T)$ will have exactly one parent $\beta \in V(T) \cup \mathcal{C}(T)$ with $\beta\alpha \in E(T)$. For a node $\alpha \in V(T) \cup \mathcal{C}(T)$ let walk(α) be the the unique path in T from a root to α. Note that walk(α) is a subgraph of T. Sometimes we will treat walks in T as sequences of nodes in $V \cup \mathcal{C}$ discussed in Section 5.1 (i.e. with timestamps removed); such places should be clear from the context.

We will grow the forest T in a greedy manner as shown in Algorithm 1. The structure of the algorithm resembles that of the Edmonds' algorithm for matchings [9], with the following important distinctions: First, in the Edmonds' algorithm each "constraint node" (i.e. each node of the input graph) can be added to the forest at most once, while in Algorithm 1 some constraints $C \in \mathcal{C}$ can be added to T and "expanded" multiple times (i.e. $E(T)$ may contain edges $C^s u$ and $C^t w$ added at distinct timestamps $s \neq t$). This is because we allow more general constraints. In particular, if C is a "perfect matching" constraint (i.e. $C = \{(a_1, \ldots, a_k) \in \{0,1\}^k \,:\, a_1 + \ldots + a_k = 1\}$) then Algorithm 1 will expand it at most once. (We will not use this fact, and thus omit the proof.)

Note that even when we enter a constraint node for the second or third time, we "branch out" based on transitions vCw available before the first visit, even though the tuple of C might have changed in the meantime. This could cause one to doubt that Algorithm 1 works at all.

A vague answer to this objection is that we grow T very carefully: While the Edmonds' algorithm does not impose any restrictions on the order in which the forest is grown, we require that all valid children $w \in W$ be added to T simultaneously when exploring edge $\{v, C\}$ in step 4. Informally speaking, this will guarantee that forest T does not have "shortcuts", a property that will be essential in the proofs. The possibility of having shortcuts is something that is not present in graph matchings and is one of the properties of even Δ-matroids responsible for the considerable length of the correctness proofs.

In the following theorem, we collect all pieces we need to show that Algorithm 1 is correct and runs in polynomial time:

Theorem 22. *If I is a CSP instance, f is a valid edge labeling of I, and we run Algorithm 1, then the following is true:*

(a) The mapping $f \oplus p$ from step 4c is a valid edge labeling of I with fewer inconsistencies than f.

(b) When contracting a blossom as described Section 5.3 I^b is an edge CSP instance with even Δ-matroid constraints and f^b is a valid edge labeling to I^b.

(c) The recursion in 4d will occur at most $O(|V|)$ many times.

Algorithm 1 Improving a given edge labeling

Input: Instance I, valid edge labeling f of I.

Output: A valid edge labeling g of I with fewer inconsistent variables than f, or "No" if no such g exists.

1. Initialize T as follows: set timestamp $t = 1$, and for each inconsistent variable $v \in V$ of I add v to T as an isolated root.
2. Pick an edge $\{v, C\} \in \mathcal{E}$ such that $v \in V(T)$ but there is no s such that $vC^s \in E(T)$ or $C^s v \in E(T)$. (If no such edge exists, then output "No" and terminate.)
3. Add new node C^t to T together with the edge vC^t.
4. Let W be the set of all variables $w \neq v$ in the scope of C such that $f(C) \oplus v \oplus w \in C$. For each $w \in W$ do the following (see Figure 2):
   (a) If $w \notin V(T)$, then add w to T together with the edge $C^t w$.
   (b) Else if w has a parent of the form C^s for some s, then do nothing.
   (c) Else if v and w belong to different trees in T (i.e. originate from different roots), then we have found an augmenting path. Let $p = \text{walk}(C^t), \text{walk}(w)^{-1}$, output $f \oplus p$ and exit.
   (d) Else if v and w belong to the same tree in T, then we have found a blossom. Form a new instance I^b and new valid edge labeling f^b of I^b by *contracting* this blossom. Solve this instance recursively, use the resulting improved edge labeling for I^b (if it exists) to compute an improved valid edge labeling for I, and terminate. All details are given in Sec. 5.3.
5. Increase the timestamp t by 1 and goto step 2.

(d) In step 4d, f^b is optimal for I^b if and only if f is an optimal for I. Moreover, given a valid edge labeling g^b of I^b with fewer inconsistent variables than f^b, we can in polynomial time output a valid edge labeling g of I with fewer inconsistent variables than f.

(e) If the algorithm answers "No" then f is optimal.

5.3 Contracting a blossom (step 4d)

We now elaborate step 4d of Algorithm 1. First, let us describe how to obtain the blossom b. Let $\alpha \in V(T) \cup \mathcal{C}(T)$ be the lowest common ancestor of nodes v and w in T. Two cases are possible.

1. $\alpha = r \in V(T)$. Variable node r must be inconsistent in f because it has outdegree two. We let $b = \text{walk}(C^t), \text{walk}(w)^{-1}$ in this case.
2. $\alpha = R^s \in \mathcal{C}(T)$. Let r be the child of R^s in T that is an ancestor of v. Replace edge labeling f with $f \oplus \text{walk}(r)$ (variable r then becomes inconsistent). Now define walk $b = p, q^{-1}, r$ where p is the walk from r to C^t in T and q is the walk from R^s to w in T (see Figure 3).

Lemma 23 (To be proved in Section A.3)**.** *Assume that Algorithm 1 reaches step 4d and one of the cases described in the above paragraph occurs. Then:*

(a) in the case 2 the edge labeling $f \oplus \mathrm{walk}(r)$ is valid, and

(b) in both cases the walk b is an f-blossom (for the new edge labeling f, in the second case). (Note that we have not formally defined f-blossoms yet; they require some machinery that will come later – see Definition 36.)

To summarize, at this point we have a valid edge labeling f of instance I and an f-blossom $b = b_0C_1b_1 \dots C_kb_k$. Let us denote by L the set of constraints in the blossom, i.e. $L = \{C_1, \dots, C_k\}$.

We construct a new instance I^b and its valid edge labeling f^b by *contracting the blossom b* as follows: we take I, add one $|L|$-ary constraint N to I, delete the variables $b_1, \dots, b_k$, and add new variables $\{v_C \colon C \in L\}$ (see Figure 4). The scope of N is $\{v_C \colon C \in L\}$ and the matroid of N consists of exactly those maps $\alpha \in \{0,1\}^L$ that send one v_C to 1 and the rest to 0.

In addition to all this, we replace each blossom constraint $D \in L$ by the constraint D^b whose scope is $\sigma \setminus \{b_1, \dots, b_k\} \cup \{v_D\}$ where σ is the scope of D. The constraint relation of D^b consists of all maps β for which there exists $\alpha \in D$ such that α agrees with β on $\sigma \setminus \{b_1, \dots, b_k\}$ and one of the following occurs (see Figure 5):

(a) $\beta(v_D) = 0$ and α agrees with $f(D)$ on all variables in $\{b_1, \dots, b_k\} \cap \sigma$, or

(b) $\beta(v_D) = 1$ and there is exactly one variable $z \in \{b_1, \dots, b_k\} \cap \sigma$ such that $\beta(z) \neq f(D)(z)$.

The lemma below follows from a more general result shown e.g. in [10, Theorem 4] (there is also a straightforward elementary proof which we include in Section A.3).

Lemma 24. *Each C_i^b is an even Δ-matroid.*

We define the edge labeling f^b of I^b as follows: for constraints $A \notin \{C_1, \dots, C_k, N\}$ we set $f^b(A) = f(A)$. For each $C \in L$, we let $f^b(C^b)(v) = f(C)(v)$ when $v \neq v_C$, and $f^b(C^b)(v_C) = 0$. Finally, we let $f^b(N)(v_C) = 1$ for $C = C_1$ and $f^b(N)(v_C) = 0$ for all other Cs. (The last choice is arbitrary; initializing $f^b(N)$ with any other tuple in N would work as well).

It is easy to check that f^b is valid for I^b. Furthermore, v_{C_1} is inconsistent in f^b while for each $C \in L \setminus \{C_1\}$ the variable v_C is consistent.

Observation 25. *In the situation described above, the instance I^b will have at most as many variables as I and one constraint more than I. Edge labelings f and f^b have the same number of inconsistent variables.*

Corollary 26 (Theorem 22(c))**.** *When given an instance I, Algorithm 1 will recursively call itself $O(|V|)$ many times.*

Proof. Since $\mathcal{C}$ and V are partitions of G_I and the degree of each $v \in V$ is two, the number of edges of G_I is $2|V|$. From the other side, the number of edges of G_I is equal to the sum of arities of all constraints in I. Since we never consider constraints with empty scopes, the number of constraints of an instance is at most double the number of variables of the instance.

Since each contraction adds one more constraint and never increases the number of variables, it follows that there can not be a sequence of consecutive contractions longer than $2|V|$, which is $O(|V|)$. □

The following two lemmas, which we prove in Appendix A, show why the procedure works. In both lemmas, we let (I, f) and (I^b, f^b) denote the instance and the valid edge labeling before and after the contraction, respectively.

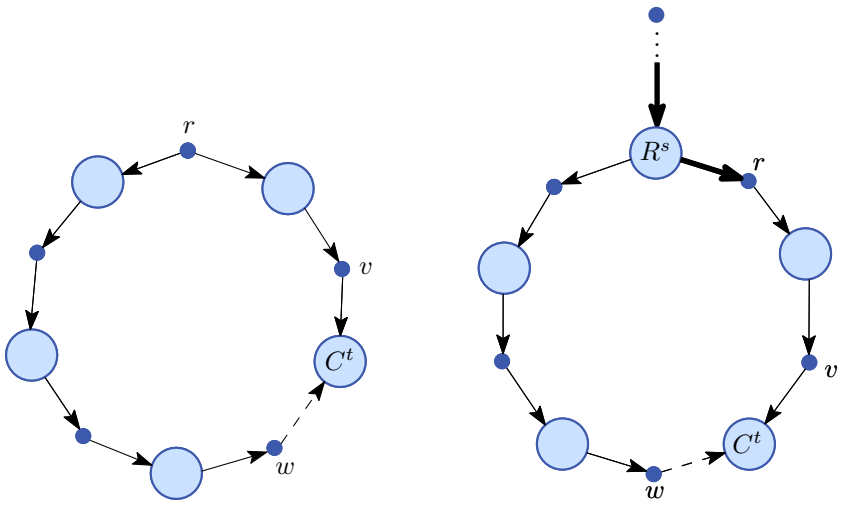

Figure 3: The two cases of step 4d. On the left, $\alpha = r$ is a variable, while on the right $\alpha = R^s$ is a constraint and the thick edges denote $p = \text{walk}(r)$. The dashed edges are orientations of edges from $\mathcal{E}$ that are not in the digraph T, but belong to the blossom.

Lemma 27. *In the situation described above, if f^b is optimal for I^b, then f is optimal for I.*

Lemma 28. *In the situation described above, if we are given a valid edge labeling g^b of I^b with fewer inconsistencies than f^b, then we can find in polynomial time a valid edge labeling g of I with fewer inconsistencies than f.*

5.4 Time complexity of Algorithm 1

To see that Algorithm 1 runs in time polynomial in the size of I, consider first the case when step 4d does happen. In this case, the algorithm runs in time polynomial in the size of I, since it essentially just searches through the graph G_I.

Moreover, from the description of contracting a blossom in part 5.3, it is easy to see that one can compute I^b and f^b from I and f in polynomial time and that I^b is not significantly larger than I: I^b has at most as many variables as I and the contracted blossom constraints C^b are not larger than the original constraints C. Finally, I^b does have one brand new constraint N, but N contains only $O(|V|)$ many tuples. Therefore, we have $|I^b| \leq |I| + O(|V|)$ where $|V|$ does not change. By Corollary 26, there will be at most $O(|V|)$ contractions in total, so the size of the final instance $I^\star$ is at most $|I| + O(|V|^2)$, which is easily polynomial in $|I|$.

All in all, Algorithm 1 will give its answer in time polynomial in $|I|$.

Appendices

A Proofs

In this appendix, we flesh out detailed proofs of the statements we gave above. In the whole section, I will be an instance of a Boolean edge CSP whose constraints are even Δ-matroids.

In Sec. A.1 we establish some properties of f-walks, and show in particular that a valid edge labeling f of I is non-optimal if and only if there exists an augmenting f-walk in I. In Sec. A.2 we

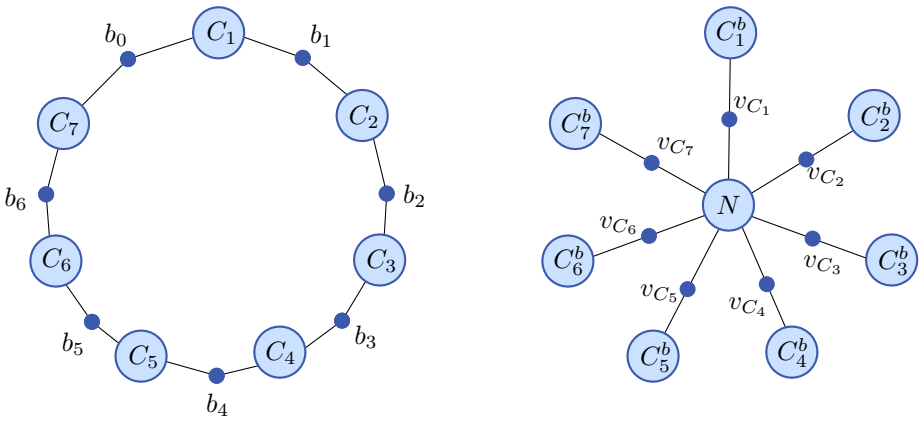

Figure 4: A blossom (left) and a contracted blossom (right) in the case when all constraints $C_1, \ldots, C_k$ are distinct. If some constraints appear in the blossom multiple times then the number of variables v_D will be smaller than k (see Figure 5).

introduce the notion of an *f-DAG*, prove that the forest T constructed during the algorithm is in fact an f-DAG, and describe some tools for manipulating f-DAGs. Then in Sec. A.3 we analyze augmentation and contraction operations, namely prove Theorem 22(a) and Lemmas 23, 27, 28 (which imply Theorem 22(b, d). Finally, in Sec. A.4 we prove Theorem 22(e).

For edge labelings f, g, let $f \,\Delta\, g \subseteq \mathcal{E}$ be the set of edges in $\mathcal{E}$ on which f and g differ.

Observation 29. *If f and g are valid edge labelings of instance I then they have the same number of inconsistencies modulo 2.*

Proof. We use induction on $|f \,\Delta\, g|$. The base case $|f \,\Delta\, g| = 0$ is trivial. For the induction step let us consider valid edge labelings f, g with $|f \,\Delta\, g| \geq 1$. Pick an edge $\{v, C\} \in f \,\Delta\, g$. By the property of even Δ-matroids there exists another edge $\{w, C\} \in f \,\Delta\, g$ with $w \neq v$ such that $f(C) \oplus v \oplus w \in C$. Thus, edge labeling $f^\star = f \oplus (vCw)$ is valid. Clearly, f and $f^\star$ have the same number of inconsistencies modulo 2. By the induction hypothesis, the same holds for edge labelings $f^\star$ and g (since $|f^\star \,\Delta\, g| = |f \,\Delta\, g| - 2$). This proves the claim. □

A.1 The properties of f-walks

Let us begin with some results on f-walks that will be of use later. The following lemma is a (bit more technical) variant of the well known property of labelings proven in [7, Theorem 3.6]:

Lemma 30. *Let f, g be valid edge labelings of I such that g has fewer inconsistencies than f, and x be an inconsistent variable in f. Then there exists an augmenting f-walk that begins in a variable different from x. Moreover, such a walk can be computed in polynomial time given I, f, g, and x.*

Proof. Our algorithm will proceed in two stages. First, we repeatedly modify the edge labeling g using the following procedure:

1. Pick a variable $v \in V$ which is consistent in f, but not in g. (If no such v exists then go to the next paragraph). By the choice of v, there exists a unique edge $\{v, C\} \in f \,\Delta\, g$. Pick a variable $w \neq v$ in the scope of C such that $\{w, C\} \in f \,\Delta\, g$ and $g(C) \oplus v \oplus w \in C$ (it exists since C is an even Δ-matroid). Replace g with $g \oplus (vCw)$, then go to the beginning and repeat.

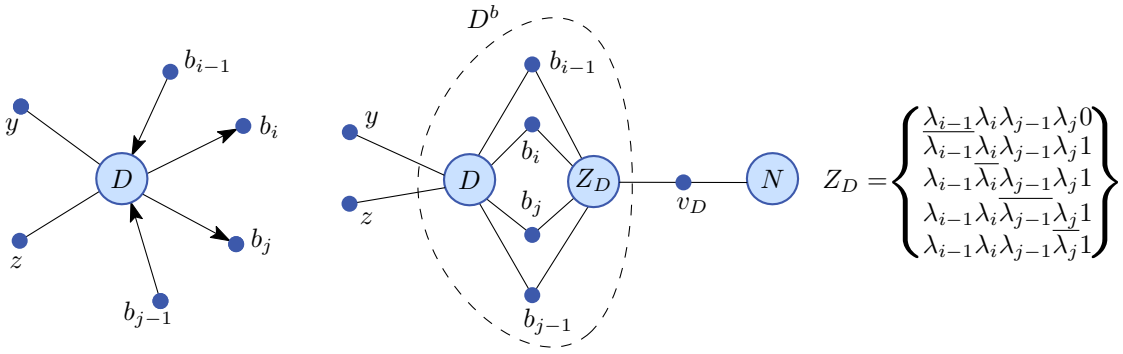

Figure 5: Modification of a constraint node D that appears in a blossom b twice, i.e. when $b = \ldots b_{i-1} D b_i \ldots b_{j-1} D b_j \ldots$ (and so $D = C_i = C_j$). Variables y and z are not part of the walk. The construction of D^b described in the text can be alternatively viewed as attaching a "gadget" constraint Z_D as shown in the figure. Here Z_D is an even Δ-matroid with five tuples that depend on the values $\lambda_k = f(\{b_k, D\})$ and $\overline{\lambda_k} = 1 - \lambda_k$.

It can be seen that g remains a valid edge labeling, and the number of inconsistencies in g never increases. Furthermore, each step decreases $|f \,\Delta\, g|$ by 2, so this procedure must terminate after at most $O(|\mathcal{E}|) = O(|V|)$ steps.

We now have valid edge labelings f, g such that f has more inconsistencies than g, and variables consistent in f are also consistent in g. Since the parity of number of inconsistencies in f and g is the same, f has at least two more inconsistent variables than g; one of them must be different from x.

In the second stage we will maintain an f-walk p and the corresponding valid edge labeling $f^\star = f \oplus p$. To initialize, pick a variable $r \in V \setminus \{x\}$ which is consistent in g but not in f, and set $p = r$ and $f^\star = f$. We then repeatedly apply the following step:

2. Let v be the endpoint of p. The variable v is consistent in g but not in $f^\star$, so there is a unique edge $\{v, C\} \in f^\star \,\Delta\, g$. Pick a variable $w \neq v$ in the scope of C such that $\{w, C\} \in f^\star \,\Delta\, g$ and $f^\star(C) \oplus v \oplus w \in C$ (it exists since C is an even Δ-matroid). Append vCw to the end of p, and accordingly replace $f^\star$ with $f^\star \oplus (vCw)$ (which is valid by the choice of w). As a result of this update, edges $\{v, C\}$ and $\{w, C\}$ are removed from $f^\star \,\Delta\, g$.

   If w is inconsistent in f, then output p (which is an augmenting f-walk) and terminate. Otherwise w is consistent in f (and thus in g) but not in $f^\star$; in this case, go to the beginning and repeat.

Each step decreases $|f^\star \,\Delta\, g|$ by 2, so this procedure must terminate after at most $O(|\mathcal{E}|) = O(|V|)$ steps. It can also be seen that p is indeed a walk. In particular, the starting node r has exactly one incident edge in the graph $(V \cup \mathcal{C}, f^\star \,\Delta\, g)$. Since this edge is immediately removed from $f^\star \,\Delta\, g$, we will never encounter the variable r again during the procedure. □

A.2 Invariants of Algorithm 1: f-DAGs

In this section we examine the properties of the forest T as generated by Algorithm 1. For future comfort, we will actually allow T to be a bit more general than what appears in Algorithm 1 – our T can be a directed acyclic digraph (DAG):

Definition 31. Let I be a Boolean edge CSP instance and f a valid edge labeling of I. We will call a directed graph T an *f-DAG* if $T = (V(T) \cup \mathcal{C}(T), E(T))$ where $V(T) \subseteq V$, $\mathcal{C}(T) \subseteq \mathcal{C} \times \mathbb{N}$, and the following conditions hold:

(a) Edges of $E(T)$ have the form vC^t or C^tv where $\{v, C\} \in \mathcal{E}$ and $t \in \mathbb{N}$.

(b) For each $\{v, C\} \in \mathcal{E}$ there is at most one $t \in \mathbb{N}$ such that vC^t or C^tv appears in $E(T)$. Moreover, vC^t and C^tv are never both in $E(T)$.

(c) Each node $v \in V(T)$ has at most one incoming edge. (Note that by the previous properties, the node v can have at most two incident edges in T.)

(d) Timestamps t for nodes $C^t \in \mathcal{C}(T)$ are all distinct (and thus give a total order on $\mathcal{C}(T)$). Moreover, this order can be extended to a total order $\prec$ on $V(T) \cup \mathcal{C}(T)$ such that $\alpha \prec \beta$ for each edge $\alpha\beta \in E(T)$. (So in particular the digraph T is acyclic.)

(e) If T contains edges uC^t and one of vC^t or C^tv, then $f(C) \oplus u \oplus v \in C$.

(f) ("No shortcuts" property) If T contains edges uC^s and one of vC^t or C^tv where $s < t$, then $f(C) \oplus u \oplus v \notin C$.

It is easy to verify that any subgraph of an f-DAG is also an f-DAG. If T is an f-DAG, then we denote by $f \oplus T$ the edge labeling we obtain from f by flipping the value of any $f(\{v, C\})$ such that $vC^t \in E(T)$ for some timestamp t. We will need to show that $f \oplus T$ is a valid edge labeling for nice enough f-DAGs T.

The following lemma shows the promised invariant property:

Lemma 32. *Let us consider the structure T during the run of Algorithm 1 with the input I and f. At any moment during the run, the forest T is an f-DAG.*

Moreover, if steps 4c or 4d are reached, then the digraph $T^\star$ obtained from T by removing all edges outgoing from C^t and adding the edge wC^t is also an f-DAG.

Proof. Obviously, an empty T is an f-DAG, as is the initial T consisting of inconsistent variables and no edges. To verify that T remains an f-DAG during the whole run of Algorithm 1, we need to make sure that neither adding vC^t in step 3, nor adding C^tw in step 4a violates the properties of T. Let us consider step 3 first. By the choice of v and C^t, we immediately get that properties (a), (b), (c), and (d) all hold even after we have added vC^t to T (we can order the nodes by the order in which they were added to T). Since there is only one edge incident with C^t, property (e) holds as well. Finally, the only way the no shortcuts property (i.e. property (f)) could fail would be if there were some u and s such that $uC^s \in E(T)$ and $f(C) \oplus u \oplus v \in C$. But then, after the node C^s got added to T, we should have computed the set W of variables w such that $f(C) \oplus v \oplus w$ (step 4) and u should have been in $W \setminus V(T)$ at that time, i.e. we should have added the edge C^sw before, a contradiction. The analysis of step 4a is similar.

Assume now that Algorithm 1 has reached one of steps 4c or 4d and consider the DAG $T^\star$ that we get from T by removing all edges of the form C^tz and adding the edge wC^t. Note that the node C^t is the only node with two incoming edges. The only three properties that this could possibly affect are (b), (e) and (f). Were (b) violated, we would have $C^sw \in E(T)$ already, and so step 4b would be triggered instead of steps 4c or 4d. For property (e), the only new pair of edges to consider is vC^t and wC^t for which we have $f(C) \oplus v \oplus w \in C$. Finally, if property (f) became violated after adding the edge wC^t then there were a u and $s < t$ such that $uC^s \in E(T)$ and $f(C) \oplus u \oplus w \in C$. Node C^s must have been added after w, or else we would have $C^sw \in E(T)$. Also, w cannot have a

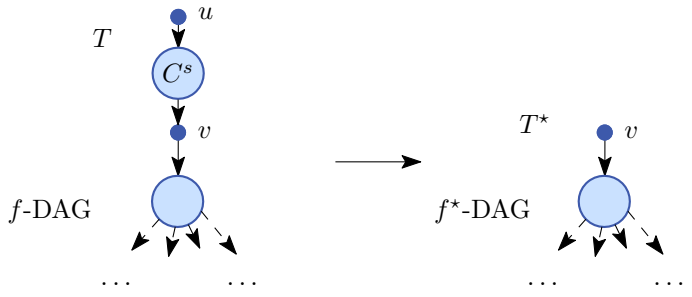

Figure 6: An f-DAG T on the left turns into $f^\star$-DAG $T^\star$ on the right; the setting from Lemma 33.

parent of the form C^k (otherwise step 4b would be triggered for w when expanding C^t). But then one of steps 4d or 4c would be triggered at timestamp s already when we tried to expand C^s, a contradiction. □

We will use the following two lemmas to prove that $f \oplus p$ is a valid edge labeling of I for various paths p that appear in steps 4c and 4d.

Lemma 33. *Let T be an f-DAG, and C^s be the constraint node in $\mathcal{C}(T)$ with the smallest timestamp s. Suppose that C^s has exactly two incident edges, namely incoming edge uC^s where u does not have other incident edges besides uC^s and another edge C^sv (see Figure 6). Let $f^\star = f \oplus (uCv)$ and let $T^\star$ be the DAG obtained from T by removing nodes u, C^s and the two edges incident to C^s. Then $f^\star$ is a valid edge labeling of I and $T^\star$ is an $f^\star$-DAG.*

Proof. Since $T^\star$ is a subgraph of T, it immediately follows that $T^\star$ satisfies the properties (a), (b), (c), and (d) from the definition of an f-DAG all hold.

Let us show that $T^\star$ has property (e). Consider a constraint node $C^t \in \mathcal{C}(T^\star)$ with $t > s$ (nothing has changed for other constraint nodes in $\mathcal{C}(T^\star)$), and suppose that $T^\star$ contains edges xC^t and one of yC^t or C^ty. If $x = y$, the situation is trivial, so assume that u, v, x, y are all distinct variables. We need to show that $f^\star(C) \oplus x \oplus y \in C$. The constraint C contains the tuples $f(C) \oplus u \oplus v$ and $f(C) \oplus x \oplus y$ (by condition (e) for T), but the no shortcuts property prohibits the tuples $f(C) \oplus u \oplus x$ and $f(C) \oplus u \oplus y$ from lying in C. Therefore, applying the even Δ-matroid property on $f(C) \oplus u \oplus v$ and $f(C) \oplus x \oplus y$ in the variable u we get that C must contain $f(C) \oplus u \oplus v \oplus x \oplus y$, so we have $f^\star(C) \oplus x \oplus y \in C$.

Now let us prove that $T^\star$ and $f^\star$ have the "no shortcuts" property. Consider constraint nodes C^k, C^ℓ in $\mathcal{C}(T^\star)$ with $s < k < \ell$ (since nothing has changed for constraint nodes other than C), and suppose that $T^\star$ contains edges xC^k and one of yC^ℓ or $C^\ell y$, where again u, v, x, y are all distinct variables. We need to show that $f^\star(C) \oplus x \oplus y \notin C$, or equivalently that $f(C) \oplus u \oplus v \oplus x \oplus y \notin C$.

Assume that it is not the case. Apply the even Δ-matroid property to tuples $f(C) \oplus u \oplus v \oplus x \oplus y$ and $f(C)$ (which are both in C) in coordinate v. We get that either $f(C) \oplus x \oplus y \in C$, or $f(C) \oplus u \oplus x \in C$, or $f(C) \oplus u \oplus y \in C$. This contradicts the "no shortcuts" property for the pair (C^k, C^ℓ), or (C^s, C^k), or (C^s, C^ℓ), respectively, and we are done. □

Corollary 34. *Let I be an edge CSP instance and f be a valid edge labeling.*

(a) Let T be an f-DAG that consists of two directed paths $x_0C_1^{t_1}x_1 \dots x_{k-1}C_k^{t_k}$ and $y_0D_1^{s_1} \dots y_{\ell-1}D_\ell^{s_\ell}$ that are disjoint everywhere except at the constraint $C_k^{t_k} = D_\ell^{s_\ell}$ (see Figure 7). Then $f \oplus T$ is a valid edge labeling of I.

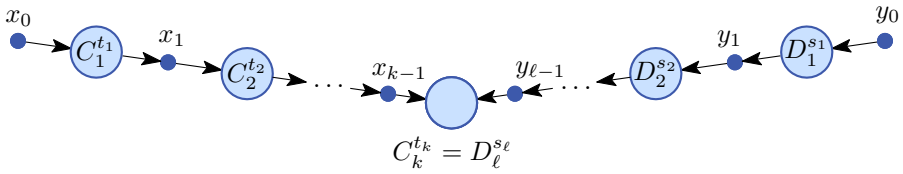

Figure 7: Two meeting paths from Corollary 34.

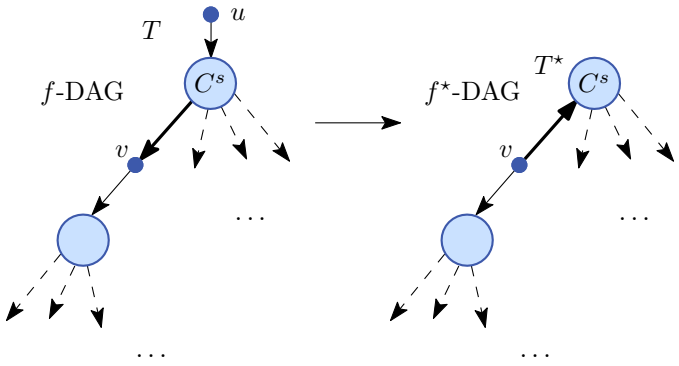

Figure 8: A more general f-DAG T turns into an $f^\star$-DAG $T^\star$ (see Lemma 35).

(b) Let T be an f-DAG that consists of a single directed path $x_0 C_1^{t_0} x_1 \dots x_{k-1} C_k^{t_k} x_k$. Then $f \oplus T$ is a valid edge labeling of I.

Proof. We will prove only part (a); the proof of part (b) is completely analogous. We proceed by induction on $k + \ell$. If $k = \ell = 1$, T consists only of the two edges $x_0 C^t$ and $y_0 C^t$ (where C^t is an abbreviated name for $C_1^{t_1} = D_1^{s_1}$). Then the fact that $f \oplus (x_0 C y_0)$ is a valid edge labeling follows from the property (e) of f-DAGs.

If we are now given an f-DAG T of the above form, then we compare t_1 and s_1. Since the situation is symmetric, we can assume without loss of generality that $s_1 > t_1$. We then use Lemma 33 for $x_1 C_1^{t_1} x_2$ (there is a x_2 since $t_k > s_1 > t_1$), obtaining the $(f \oplus (x_1 C_1 x_2))$-DAG $T^\star$ that consists of two directed paths $x_2 \dots x_k C^{t_k}$ and $y_1 D_1^{s_1} \dots y_\ell D_\ell^{s_\ell}$. Since $T^\star$ is shorter than T, the induction hypothesis gets us that $f \oplus (x_1 C_1 x_2) \oplus T^\star = f \oplus T$ is a valid edge labeling. □

Lemma 35. *Let T be an f-DAG, and C^s be the constraint node in $\mathcal{C}(T)$ with the smallest timestamp s. Suppose that C^s has exactly one incoming edge uC^s, and u does not have other incident edges besides uC^s. Suppose also that C^s has an outgoing edge $C^s v$. Let $f^\star = f \oplus (uCv)$, and $T^\star$ be the DAG obtained from T by removing the edge uC^s together with u and reversing the orientation of edge $C^s v$ (see Figure 8). Then $f^\star$ is a valid edge labeling of I and $T^\star$ is an $f^\star$-DAG.*

Proof. It is easy to verify that $T^\star$ satisfies the properties (a), (b) and (c). To see property (d), just take the linear order on nodes of T and change the position of v so that it is the new minimal element in this order (v has no incoming edges in $T^\star$).

Let us prove that property (e) of Definition 31 is preserved. First, consider constraint node C^s. Suppose that $T^\star$ contains one of xC^s or $C^s x$ with $x \neq v$. We need to show that $f^\star(C) \oplus v \oplus x \in C$, or equivalently $f(C) \oplus u \oplus x \in C$ (since $f^\star(C) \oplus v = f(C) \oplus (u \oplus v) \oplus v = f(C) \oplus u$). This claim holds by property (e) of Definition 31 for T.

Now consider a constraint node $C^t \in \mathcal{C}(T^\star)$ with $t > s$, and suppose that $T^\star$ contains edges xC^t and one of yC^t or $C^t y$. We need to show that $f^\star(C) \oplus x \oplus y \in C$, or equivalently that $f(C) \oplus u \oplus v \oplus x \oplus y \in C$. For that we can simply repeat word-by-word the argument used in the proof of Lemma 33.

Now let us prove that the "no shortcuts" property is preserved. First, consider a constraint node C^t in $\mathcal{C}(T^\star)$ with $t > s$, and suppose that $T^\star$ contains one of xC^t or $C^t x$. We need to show that $f^\star(C) \oplus v \oplus x \notin C$, or equivalently $f(C) \oplus u \oplus x \notin C$. This claim holds by the "no shortcuts" property for T. Now consider constraint nodes C^k, C^ℓ in $\mathcal{C}(T^\star)$ with $s < k < \ell$, and suppose that $T^\star$ contains edges xC^k and one of yC^ℓ or $C^\ell y$. Note that u, v, x, y are all distinct variables. We need to show that $f^\star(C) \oplus x \oplus y \notin C$, or equivalently that $f(C) \oplus u \oplus v \oplus x \oplus y \notin C$. For that we can simply repeat word-by-word the argument used to show the no shortcuts property in the proof of Lemma 33. □

A.3 Analysis of augmentations and contractions

First, we prove the correctness of the augmentation operation, i.e. that edge labeling $f \oplus p$ in step 4c is valid.

Proof. (Proof of Theorem 22(a)) Let T_1 be the f-DAG constructed during the run of Algorithm 1; let T_2 be the DAG obtained from T_1 by adding the edge wC^t. By Lemma 32, T_2 is an f-DAG. Let T_3 be the subgraph of T_2 induced by the nodes in p. It is easy to verify that T_3 consists of two directed paths that share their last node. Therefore, by Corollary 34, we get that $f \oplus T_3 = f \oplus p$ is a valid edge labeling of I. □

In the remainder of this section we show the correctness of the contraction operation by proving Lemmas 23, 27, 28. Let us begin by giving a full definition of a blossom:

Definition 36. Let f be a valid edge labeling. An *f-blossom* is any walk $b = b_0 C_1 b_1 C_2 \dots C_k b_k$ with $b_0 = b_k$ such that:

(a) variable $b_0 = b_k$ is inconsistent in f while variables $b_1, \dots, b_{k-1}$ are consistent, and

(b) there exists $\ell \in [1, k]$ and timestamps $t_1, \dots, t_k$ such that the DAG consisting of two directed paths $b_0 C_1^{t_1} \dots b_{\ell-1} C_\ell$ and $b_k C_k^{t_k} b_{k-1} \dots b_\ell C_\ell^{t_\ell}$ is an f-DAG.

Lemma 37. *Let b be a blossom. Then $b_{[i,j]}$ is an f-walk for any non-empty proper subinterval $[i, j] \subsetneq [0, k]$.*

Proof. Let us denote the f-DAG from the definition of a blossom by B. By taking an appropriate subgraph of B and applying Corollary 34 we get that $f \oplus b_{[i,j]}$ is valid for any non-empty subinterval $[i, j] \subsetneq [0, k]$. Since the set of these intervals is downward closed, $b_{[i,j]}$ is in fact an f-walk. □

Lemma (Lemma 23 restated)**.** *Assume that Algorithm 1 reaches step 4d and one of the cases described at the beginning of Section 5.3 occurs. Then:*

(a) in the case 2 the edge labeling $f \oplus \mathrm{walk}(r)$ is valid, and

(b) in both cases the walk b is an f-blossom (for the new edge labeling f, in the second case).

Proof. Let T be the forest at the moment of contraction, $T^\dagger$ be the subgraph of T containing only paths walk(C^t) and walk(w), and $T^\star$ be the graph obtained from $T^\dagger$ by adding the edge wC^t. By Lemma 32 graph $T^\star$ is an f-DAG (we also need to observe that any subgraph of an f-DAG is again an f-DAG).

If the lowest common ancestor of u and v in T is a variable node $r \in V(T)$ (i.e. we have case 1 from Section 5.3), then the f-DAG $T^\star$ consists of two directed paths from r to the constraint C and it is easy to verify that when we let b to be one of these paths followed by the other in reverse, we get a blossom.

Now consider case 2, i.e. when the lowest common ancestor of u and v in T is a constraint node $R^s \in \mathcal{C}(T)$. Note that $T^\star$ has the unique source node u (that does not have incoming edges), and u has an outgoing edge uD^t where D^t is the constraint node with the smallest timestamp in $T^\star$. Let us repeat the following operation while $D^t \neq R^s$: Replace f with $f \oplus (uD^t z)$ where z is the unique out-neighbor of D^t in $T^\star$, and simultaneously modify $T^\star$ by removing nodes u, D^t and edges $uD^t, D^t z$. By Lemma 33 f remains a valid edge labeling throughout this process, and $T^\star$ remains an f-DAG (for the latest f).

We get to the point that the unique in-neighbor u of R^s is the source node of $T^\star$. Replace f with $f \oplus (uR^s r)$, and simultaneously modify $T^\star$ by removing node u together with the edge uR^s and reversing the orientation of edge $R^s r$. The new f is again valid, and the new $T^\star$ is an f-DAG by Lemma 35. This means that the resulting walk b is an f-blossom for the new f. □

Lemma (Lemma 24 restated)**.** *Each D^b constructed in Section 5.3 is an even Δ-matroid.*

Proof. We will denote by σ the scope of D. Let $\beta, \beta' \in D^b$ and assume that $\beta(v) \neq \beta'(v)$ for some v. We will show that there is a $u \neq v$ such that $\beta(u) \neq \beta'(u)$ and $\beta \oplus v \oplus u \in D^b$.

Assume first that v is different from v_D. From the definition of D^b, we know that there exist $\alpha, \alpha' \in D$ that witness $\beta, \beta' \in D^b$. In particular, $\alpha(v) \neq \alpha'(v)$. Since D is an even Δ-matroid, there exists a variable w such that $\alpha(w) \neq \alpha'(w)$ and $\alpha \oplus v \oplus w \in D$. If w is not in $\{b_1, \ldots, b_k\}$, we get $\beta(w) \neq \beta'(w)$ and $\beta \oplus v \oplus w \in D^b$, while if $w \in \{b_1, \ldots, b_k\}$, we get $\beta \oplus v \oplus v_D \in D^b$. If $\beta(v_D) \neq \beta'(v_D)$, we are done. Otherwise, α and α' disagree exactly on two variables from $\{b_1, \ldots, b_k\}$, one of which is w; let t be the other variable. The Δ-matroid property used on $\alpha \oplus v \oplus w$ and α' gives us that there is a $z \notin \{b_1, \ldots, b_k\}$ such that $\alpha(z) \neq \alpha'(z)$ and $\alpha \oplus v \oplus w \oplus t \oplus z \in D$. From this, it easily follows that $\beta \oplus v \oplus z \in D^b$.

In the case $v = v_D$, we again consider $\alpha, \alpha' \in D$ as above. Since $\beta(v_D) \neq \beta'(v_D)$, we know that α and α' disagree at exactly one variable b_i. Using the Δ-matroid property on α, α' with the variable b_i, we obtain a variable $w \neq b_i$ such that α and α' disagree on w and $\alpha \oplus b_i \oplus w \in D$. Then w needs to be from $\sigma \setminus \{b_1, \ldots, b_k\}$, so $\beta(w) \neq \beta'(w)$ and $\alpha \oplus b_i \oplus w$ witnesses that $\beta \oplus v_D \oplus w \in D^b$ □

Finally, we prove two lemmas showing that if we contract a blossom b in instance I to obtain the instance I^b and the edge labeling f^b, then f^b is optimal if and only if f^b is optimal for I^b.

Lemma (Lemma 27 restated)**.** *In the situation described above, if f^b is optimal for I^b, then f is optimal for I.*

Proof. Assume that f is not optimal for I, so there exists a valid edge labeling g with fewer inconsistencies than f. Then by Lemma 30 there exists an augmenting f-walk p in I that starts at some node other than b_k. Denote by p^b the sequence obtained from p by replacing each C_i from the blossom by C_i^b. Observe that if p does not contain the variables $b_1, \ldots, b_k$, then p is an f-walk if and only if p^b is an f^b-walk, so the only interesting case is when p enters the set $\{b_1, \ldots, b_k\}$.

We will proceed along p and consider the first i such that there is a blossom constraint D and an index j for which $p_{[0,i]} D b_j$ is an f-walk (i.e. we can enter the blossom from p).

If $D = C_1$, then it follows from the definition of C_1^b that $p_{[0,i]}C_1^b v_{C_1}$ is an augmenting f^b-walk in I^b, while if $D \neq C_1$, then $p_{[0,i]}D^b v_{C_D} N v_{C_1}$ is an augmenting f^b-walk. In both cases, we get an augmenting walk for f^b, and so f^b was not optimal. □

To show the other direction, we will first prove the following result.

Lemma 38. *Let q be an f-walk and T an f-DAG such that there is no proper prefix $q^\star$ of q and no edge vC^s or $C^s v$ of T such that $q^\star Cv$ would be an f-walk. Then T is a $(f \oplus q)$-DAG.*

Proof. We proceed by induction on the length of q. If q has length 0, the claim is trivial. Otherwise, let $q = xCyq^\dagger$ for some $q^\dagger$. Note that $q^\dagger$ is trivially an $(f \oplus (xCy))$-walk. We verify that T is an $(f \oplus (xCy))$-DAG, at which point it is straightforward to apply the induction hypothesis with $f \oplus xCy$ and $q^\dagger$ to show that T is an $(f \oplus q)$-DAG.

We choose the timestamp t to be smaller than any of the timestamps appearing in T and construct the DAG $T^\dagger$ from T by adding the nodes x, y, C^t and edges xC^t and $C^t y$. It is easy to see that $T^\dagger$ is an f-DAG – the only property that might possibly fail is the no shortcuts property. However, since the timestamp of C^t is minimal, were the no shortcuts property violated, T would have to contain an edge of the form vC^s or $C^s v$ such that $f(C) \oplus x \oplus v \in C$. But in that case, we would have the f-walk xCv, contradicting our assumption on prefixes of q.

It follows that $T^\dagger$ is an f-DAG and we can use Lemma 33 with the constraint C^t and edges xC^t and $C^t y$ to show that T is an $(f \oplus (xC^t y))$-DAG, concluding the proof. □

Lemma (Lemma 28 restated)**.** *In the situation described above, if we are given a valid edge labeling g^b of I^b with fewer inconsistencies than f^b, then we can find in polynomial time a valid edge labeling g of I with fewer inconsistencies than f.*

Proof. Our overall strategy here is to take an inconsistency from the outside of the blossom b and bring it into the blossom. We begin by showing how to get a valid edge labeling f' for I with an inconsistent variable just one edge away from b.

Using Lemma 30, we can use g^b and f^b to find in polynomial time an augmenting f^b-walk p^b that does not begin at the inconsistent variable v_{C_1}. If p^b does not contain any of the variables $v_{C_1}, \ldots, v_{C_k}$, then we can just output the walk p obtained from p^b by replacing each C_i^b by C_i and be done. Assume now that some v_C appears in p^b. We choose the f^b-walk r^b so that $r^b C^b v_C$ is the shortest prefix of p^b that ends with some blossom variable v_C. By renaming all C^bs in r^b to Cs, we get the walk r. It is straightforward to verify that r is an f-walk and that $rC_i b_i$ or $rC_i b_{i-1}$ is an f-walk for some $i \in [1, k]$. Let q be the shortest prefix of r such that one of $qC_i b_i$ or $qC_i b_{i-1}$ is an f-walk for some $i \in [1, k]$.

Recall that the blossom b originates from an f-DAG B. The minimality of q allows us to apply Lemma 38 and obtain that B is also an $(f \oplus q)$-DAG. Let $f' = f \oplus q$ and let x be the last variable in q. It is easy to see that f' is a valid edge labeling with exactly as many inconsistent variables as f. Moreover x is inconsistent in f' and there is an index i such that at least one of $xC_i b_i$ or $xC_i b_{i-1}$ is an f'-walk. We will now show how to improve f'.

If the constraint C_i appears only once in the blossom b, it is easy to verify (using Lemma 37) that one of $xC_i b_{[i,k]}$ or $xC_i b_{[0,i-1]}{}^{-1}$ is an augmenting f'-walk. However, since the constraint C_i might appear in the blossom several times, we have to come up with a more elaborate scheme to handle that. The blossom b comes from an f'-DAG B in which some node $C_\ell^{t_\ell}$ is the node with the maximal timestamp (for a suitable $\ell \in [1, k]$). Assume first that there is a $j \in [\ell, k]$ such that $xC_j b_j$ is an f'-walk. In that case, we take maximal such j and consider the DAG B' we get by adding the edge $C_j^{t_j} x$ to the subgraph of B induced by the nodes $C_j^{t_j}, b_j, C_{j+1}^{t_{j+1}}, \ldots, C_k^{t_k}, b_k$.

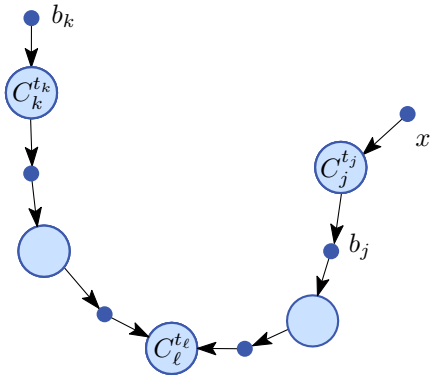

Figure 9: The f'-DAG B' constructed using xC_jb_j (where $j < \ell$).

It is routine to verify that B' is an f'-DAG; the only thing that could possibly fail is the no shortcuts property involving $C_j^{t_j}$. However, $C_j^{t_j}$ has maximal timestamp in B' and there is no $i > j$ such that $f'(C_j) \oplus x \oplus b_i \in C_j$.

Using Corollary 34, we get that $f' \oplus B'$ is a valid edge labeling which has fewer inconsistencies than f', so we are done. In a similar way, we can improve f' when there exists a $j \in [1, \ell]$ such that xC_jb_{j-1} is an f'-walk.

If neither of the above cases occurs, then we take j such that the timestamp t_j is maximal and either xC_jb_j or xC_jb_{j-1} is an f'-walk. Without loss of generality, let xC_jb_j be an f'-walk. Then $j < \ell$ and we consider the DAG B' we get from the subgraph of B induced by $C_j^{t_j}, b_j, C_{j+1}^{t_{j+1}}, \ldots, C_k^{t_k}, b_k$ by adding the edge $xC_j^{t_j}$ (see Figure 9). As before, the only way B' can not be an f'-DAG is if the no shortcuts property fails, but that is impossible: we chose j so that t_j is maximal, so an examination of the makeup of B' shows that the only bad thing that could possibly happen is if there were an index $i \geq \ell$ such that $C_i = C_j$, we had in B the edge $b_iC_i^{t_i}$, and $f'(C_i) \oplus b_i \oplus x \in C_i$. But then we would have the f'-walk xC_ib_i for $i \geq \ell$ and the procedure from the previous paragraph would apply. Using Corollary 34, we again see that $f' \oplus B'$ is a valid edge labeling with fewer inconsistencies than f.

It is easy to verify that finding q, calculating $f' = f \oplus q$, finding an appropriate j and augmenting f' can all be done in time polynomial in the size of the instance. □

A.4 Proof of Theorem 22(e)

In this section we will prove that if the algorithm answers "No" then f is an optimal edge labeling.

Lemma 39. *Suppose that Algorithm 1 outputs "No" in step 2, without ever visiting steps 4c and 4d. Then f is optimal.*

Proof. Let T be the forest upon termination, and denote

$$\begin{aligned}\overline{E}(T) =& \{Cv \mid C^t v \in E(T) \text{ for some } t\} \cup \\ & \{vC \mid vC^t \in E(T) \text{ for some } t\}.\end{aligned}$$

Inspecting Algorithm 1, one can check that $\overline{E}(T)$ has the following properties:

(a) If v is an inconsistent variable in f and $\{v, C\} \in \mathcal{E}$, then $vC \in \overline{E}(T)$.

(b) If $Cv \in \overline{E}(T)$ and $\{v, D\} \in \mathcal{E}, D \neq C$, then $vD \in \overline{E}(T)$.

(c) If $vC \in \overline{E}(T)$, then $Cv \notin \overline{E}(T)$.

(d) Suppose that $vC \in \overline{E}(T)$ and $f(C) \oplus v \oplus w \in C$ where v, w are distinct nodes in the scope of constraint C. Then $Cw \in \overline{E}(T)$.

An f-walk p will be called *bad* if it starts at a variable node which is inconsistent in f, and contains an edge $Cv \notin \overline{E}(T)$; otherwise p is *good*. Clearly, any augmenting f-walk is bad: its last edge Cv satisfies $vC \in \overline{E}(T)$ by property a, and thus $Cv \notin \overline{E}(T)$. Thus, if f is not optimal, then there exists at least one bad f-walk. Let p be a shortest bad f-walk. Write $p = p^\star(vCw)$ where $p^\star$ ends at v. By minimality of p, $p^\star$ is good and $Cw \notin \overline{E}(T)$. Using properties (a) or (b), we obtain that $vC \in \overline{E}(T)$ (and therefore $Cv \notin \overline{E}(T)$).

Let q be the shortest prefix of $p^\star$ (also an f-walk) such that $f \oplus q \oplus (vCw)$ is valid (at least one such prefix exists, namely $q = p^\star$). The walk q must be of positive length (otherwise the precondition of property (d) would hold, and we would get $Cw \in \overline{E}(T)$, a contradiction). Also, the last constraint node in q must be C, otherwise we could have taken a shorter prefix. Thus, we can write $q = q^\star(xCy)$ where $q^\star$ ends at x. Note that, since p is a walk, the variables x, y, v, w are (pairwise) distinct.

We shall write $g = f \oplus q^\star$. Let us apply the even Δ-matroid property to tuples $g(C) \oplus x \oplus y \oplus v \oplus w$ and $g(C)$ (which are both in C) in coordinate y. We get that either $g(C) \oplus v \oplus w \in C$, or $g(C) \oplus x \oplus v \in C$, or $g(C) \oplus x \oplus w \in C$. In the first case we could have chosen $q^\star$ instead of q – a contradiction to the minimality of q. In the other two cases $q^\star(xCu)$ is an f-walk for some $u \in \{v, w\}$. But then from $Cu \notin \overline{E}(T)$ we get that $q^\star(xCu)$ is a bad walk – a contradiction to the minimality of p. □

Corollary 40 (Theorem 22(e))**.** *If Algorithm 1 answers "No", then the edge labeling f is optimal.*

Proof. Algorithm 1 can answer "No" for two reasons: either the forest T can not be grown further and neither an augmenting path nor a blossom are found, or the algorithm finds a blossom b, contracts it and then concludes that f^b is optimal for I^b. We proceed by induction on the number of contractions that have occurred during the run of the algorithm.

The base case, when there were no contractions, follows from Lemma 39. The induction step is an easy consequence of Lemma 27: If we find b and the algorithm answers "No" when run on f^b and I^b then, by the induction hypothesis, f^b is optimal for I^b, and by Lemma 27 f is optimal for I. □

B Non matching realizable even Δ-matroid

Here we prove Proposition 17 which says that not every even Δ-matroid of arity six is matching realizable. We do it by first showing that matching realizable even Δ-matroids satisfy certain decomposition property and then we exhibit an even Δ-matroid of arity six which does not possess this property and thus is not matching realizable.

Lemma 41. *Let M be a matching realizable even Δ-matroid and let $f, g \in M$. Then $f \,\Delta\, g$ can be partitioned into pairs of variables $P_1, \dots P_k$ such that $f \oplus P_i \in M$ and $g \oplus P_i \in M$ for every $i = 1 \dots k$.*

Proof. Fix a graph $G = (N, E)$ that realizes M and let $V = \{v_1, \dots, v_n\} \subseteq N$ be the nodes corresponding to variables of M. Let E_f and E_g be the edge sets from matchings that correspond

to tuples f and g. Now consider the graph $G' = (N, E_f \,\Delta\, E_g)$ (symmetric difference of matchings). Since both E_f and E_g cover each node of $N \setminus V$, the degree of all such nodes in G' will be zero or two. Similarly, the degrees of nodes in $(V \setminus (f \,\Delta\, g))$ are either zero or two leaving $f \,\Delta\, g$ as the set of nodes of odd degree, namely of degree one. Thus G' is a union of induced cycles and paths, where the paths pair up the nodes in $f \,\Delta\, g$. Let us use this pairing as $P_1, \ldots, P_k$.

Each such path is a subset of E and induces an alternating path with respect to both E_f and E_g. After altering the matchings accordingly, we obtain new matchings that witness $f \oplus P_i \in M$ and $g \oplus P_i \in M$ for every i. □

Lemma 42. *There is an even Δ-matroid of arity 6 which does not have the property from Lemma 41.*

Proof. Let us consider the set M with the following tuples:

000000	100100	011011	111111
	011000	100111	
	001100	110011	
	001010	110101	
	000101	111010	
	001001	001111	
	010001	101101	
	100010	101011	
		111100	

With enough patience or with computer aid one can verify that this is indeed an even Δ-matroid. However, there is no pairing satisfying the conclusion of Lemma 41 for tuples $f = 000000$, and $g = 111111$. In fact the set of pairs P for which both $f \oplus P \in M$ and $g \oplus P \in M$ is $\{v_1, v_4\}$, $\{v_2, v_3\}$, $\{v_3, v_4\}$, $\{v_3, v_5\}$, $\{v_4, v_6\}$ (see the first five lines in the middle of the table above) but no three of these form a partition on $\{v_1, \ldots, v_6\}$. □

C Extending our algorithm to effectively coverable Δ-matroids

In this section we extend Algorithm 1 from even Δ-matroids to a wider class of so-called effectively coverable Δ-matroids. The idea of the algorithm is similar to what [11] previously did for $\mathcal{C}$-zebra Δ-matroids, but our method covers a larger class of Δ-matroids.

Let us begin by giving a formal definition of effectively coverable Δ-matroids.

Definition 43. We say that a class of Δ-matroids Γ is *effectively coverable* if there is an algorithm that, given input $M \in \Gamma$ and $\alpha \in M$, finds and lists in a polynomial time a set M_α that satisfies the conditions of Definition 8.

Before we go on, we would like to note that coverable Δ-matroids are closed under gadgets, i.e. the "supernodes" shown in Figure 1. Taking gadgets is a common construction in the CSP world, so being closed under gadgets makes coverable Δ-matroids a very natural class to study.

Definition 44. Given $M \subseteq \{0,1\}^U$ and $N \subseteq \{0,1\}^V$ with U, V disjoint sets of variables, we define the *direct product* of M and N as

$$M \times N = \{(\alpha, \beta) \colon \alpha \in M, \beta \in N\} \subseteq \{0,1\}^{U \cup V}.$$

If $w_1, w_2 \in U$ are distinct variables of M, then the Δ-matroid obtained from M by identifying w_1 and w_2 is

$$\begin{aligned} M_{w_1=w_2} &= \{\beta_{\restriction U\setminus\{w_1,w_2\}} \colon \beta \in M,\, \beta(w_1) = \beta(w_2)\} \\ &\subseteq \{0,1\}^{U\setminus\{w_1,w_2\}} \end{aligned}$$

It is easy to verify that a direct product of two (even) Δ-matroids is an (even) Δ-matroid. Showing that identifying two variables in an (even) Δ-matroid yields an (even) Δ-matroid is a bit more difficult, but still straightforward.

If a matroid P is obtained from some matroids $M_1, \ldots, M_k$ by a sequence of direct products and identifying variables, we say that P is *gadget-constructed* from $M_1, \ldots, M_k$ (a gadget is an edge CSP instance with some variables present in only one constraint – these are the "output variables").

Theorem 45. *The class of coverable even Δ-matroids is closed under:*

a) Direct products,

b) identifying pairs of variables, and

c) gadget constructions.

Proof. a) Let $M \subseteq \{0,1\}^U$ and $N \subseteq \{0,1\}^V$ be two coverable Δ-matroids. We claim that if (α, β) and (γ, δ) are even-neighbors in $M \times N$, then either $\alpha = \gamma$ and β is an even-neighbor of δ in N, or α is an even-neighbor of γ in M and $\beta = \delta$. This is straightforward to verify: Without loss of generality let us assume that $u \in U$ is a variable of M such that $(\alpha, \beta) \oplus u \notin M \times N$, and let v be the variable such that $(\alpha, \beta) \oplus u \oplus v = (\gamma, \delta)$. Since we are dealing with a direct product, we must have $\alpha \oplus u \notin M$ and in order for $(\alpha, \beta) \oplus u \oplus v$ to lie in $M \times N$, we must have $\alpha \oplus u \oplus v \in M$. But then $\delta = \beta$ and $\alpha \oplus u \oplus v = \gamma$ is an even-neighbor of α.

Let $\alpha \in M$, $\beta \in N$ and let M_α, N_β be the even Δ-matroids from Definition 8 for M and N. From the above paragraph, it follows by induction that whenever (γ, δ) is reachable from (α, β), then $(\gamma, \delta) \in M_\alpha \times N_\beta$. Since each $M_\alpha \times N_\beta$ is an even Δ-matroid, the direct product $M \times N$ satisfies the first two parts of Definition 8.

It remains to show that if we can reach $(\gamma, \delta) \in M \times N$ from $(\alpha, \beta) \in M \times N$ and $(\gamma, \delta) \oplus u \oplus v \in M_\alpha \times N_\beta \setminus M \times N$, then $(\gamma, \delta) \oplus u, (\gamma, \delta) \oplus v \in M \times N$. By the first paragraph of this proof, we can reach γ from α in M and δ from β in N. Moreover, both u and v must lie in the same set U or V, for otherwise we would have that $(\gamma \oplus v, \delta \oplus u)$ or $(\gamma \oplus u, \delta \oplus v)$ lies in $M_\alpha \times N_\beta$, a contradiction with M_α being an even Δ-matroid. So let (again without loss of generality) $u, v \in V$. Then $\delta \oplus u \oplus v \in N_\beta \setminus N$. Since N is coverable and δ is reachable from β, we get $\delta \oplus v, \delta \oplus u \in N$, giving us $(\gamma, \delta) \oplus u, (\gamma, \delta) \oplus v \in M \times N$ and we are done.

b) Let $M \subseteq \{0,1\}^U$ be coverable and $w_1 \neq w_2$ be two variables.

Similarly to the previous item, the key part of the proof is to show that the relation of being reachable survives identifying w_1 and w_2: More precisely, take $\alpha, \gamma \in M_{w_1=w_2}$ and $\beta \in M$ such that $\beta(w_1) = \beta(w_2)$ and $\alpha = \beta_{\restriction U\setminus\{w_1,w_2\}}$ (i.e. β witnesses $\alpha \in M_{w_1=w_2}$). Assume that we can reach γ from α. Then we can reach from β a tuple $\delta \in M$ such that $\delta(w_1) = \delta(w_2)$ and $\gamma = \delta_{\restriction U\setminus\{w_1,w_2\}}$.

Since we can proceed by induction, it is enough to prove this claim in the case when α, γ are even-neighbors. So assume that there exist variables u and v such that $\gamma = \alpha \oplus u \oplus v$ and

$\alpha \oplus u \notin M_{w_1=w_2}$. From the latter, it follows that $\beta \oplus u, \beta \oplus u \oplus w_1 \oplus w_2 \notin M$. Knowing all this, we see that if $\beta \oplus u \oplus v \in M$, the tuple $\beta \oplus u \oplus v$ is an even-neighbor of β and we are done.

Assume thus that $\beta \oplus u \oplus v \notin M$. Let δ be the tuple of M witnessing $\gamma \in M_{w_1=w_2}$. Since $\beta \oplus u \oplus v \notin M$, we get $\delta = \beta \oplus u \oplus v \oplus w_1 \oplus w_2$. Since $\beta \oplus u, \beta \oplus u \oplus v \notin M$, the Δ-matroid property applied on β and δ in the variable u gives us (without loss of generality) that $\beta \oplus u \oplus w_1 \in M$. But then β is an even-neighbor of $\beta \oplus u \oplus w_1$ in M, which is an even neighbor (via the variable w_2 – recall that $\beta \oplus u \oplus w_1 \oplus w_2 \notin M$) of δ and so we can reach δ from β, proving the claim.

Assume now that M is coverable. We want to show that the sets $(M_\beta)_{w_1=w_2}$ where β ranges over M cover $M_{w_1=w_2}$. Choose $\alpha \in M_{w_1=w_2}$ and let $\beta \in M$ be the witness for $\alpha \in M_{w_1=w_2}$. We claim that the even Δ-matroid $(M_\beta)_{w_1=w_2}$ contains all members of $M_{w_1=w_2}$ that can be reached from α. Indeed, whenever γ can be reached from α, some $\delta \in M$ that witnesses $\gamma \in M_{w_1=w_2}$ can be reached from β, so $\delta \in M_\beta$ and $\gamma \in (M_\beta)_{w_1=w_2}$.

To finish the proof, take $\beta \in M$ witnessing $\alpha \in M_{w_1=w_2}$ and $\gamma \in M_{w_1=w_2}$ that is reachable from α and satisfies $\gamma \oplus u \oplus v \in (M_\beta)_{w_1=w_2} \setminus M_{w_1=w_2}$ for a suitable pair of variables u, v. Take a $\delta \in M$ that witnesses $\gamma \in M_{w_1=w_2}$ and is reachable from β (we have shown above that such a δ exists). Since $\gamma \oplus u \oplus v \in (M_\beta)_{w_1=w_2} \setminus M_{w_1=w_2}$, we know that neither $\delta \oplus u \oplus v$ nor $\delta \oplus u \oplus v \oplus w_1 \oplus w_2$ lies in M, but at least one of these two tuples lies in M_β. If $\delta \oplus u \oplus v \in M_\beta$, we just use coverability of M to get $\delta \oplus u, \delta \oplus v \in M$, which translates to $\gamma \oplus u, \gamma \oplus v \in M_{w_1=w_2}$. If this is not the case, we know that $\delta, \delta \oplus u \oplus v \oplus w_1 \oplus w_2 \in M_\beta$ and $\delta \oplus u \oplus v \notin M_\beta$. We show that in this situation we have $\gamma \oplus u \in M_{w_1=w_2}$; the proof of $\gamma \oplus v \in M_{w_1=w_2}$ is analogous.

Using the even Δ-matroid property of M_β on δ and $\delta \oplus u \oplus v \oplus w_1 \oplus w_2$ in the variable u, we get that without loss of generality $\delta \oplus u \oplus w_1 \in M_\beta$ (recall that $\delta \oplus u \oplus v \notin M_\beta$). If $\delta \oplus u \oplus w_1 \notin M$, we can directly use coverability of M on δ to get that $\delta \oplus u \in M$, resulting in $\gamma \oplus u \in M_{w_1=w_2}$. If, on the other hand, $\delta \oplus u \oplus w_1 \in M$ and $\delta \oplus u \notin M$, then $\delta \oplus u \oplus w_1$ is reachable from β, so we can use coverability of M on $\delta \oplus u \oplus w_1 \in M$ and $\delta \oplus u \oplus v \oplus w_1 \oplus w_2 \in M_\beta \setminus M$ to get $\delta \oplus u \oplus w_1 \oplus w_2 \in M$, which again results in $\gamma \oplus u \in M_{w_1=w_2}$, finishing the proof.

c) This follows from the previous two points as any gadget construction is equivalent to a sequence of products followed by identifying variables.

□

Going back to edge CSP, we can reuse many notions from the even Δ-matroid case. The definitions of valid, optimal, and non-optimal edge labeling may remain intact for coverable Δ-matroids, but we need to adjust our definition of a walk, which will now be allowed to end in a constraint.

Definition 46 (Walk for general Δ-matroids)**.** A *walk* q of length k or $k + 1/2$ in the instance I is a sequence $q_0 C_1 q_1 C_2 \dots C_k q_k$ or $q_0 C_1 q_1 C_2 \dots C_{k+1}$, respectively, where the variables q_{i-1}, q_i lie in the scope of the constraint C_i, and each edge $\{v, C\} \in \mathcal{E}$ is traversed at most once: vC and Cv occur in q at most once, and they do not occur simultaneously.

Given an edge labeling f and a walk q, we define the edge labeling $f \oplus q$ in the same way as before (see eq. (1)). We also extend, in a natural way, the definitions of an f-walk and an augmenting f-walk for a valid edge labeling f: A walk q is an f-walk if $f \oplus q^\star$ is a valid edge labeling whenever $q^\star = q$ or $q^\star$ is a prefix of q that ends at a variable. An f-walk is called augmenting if: (1) it starts at a variable inconsistent in f, (2) it ends either at a different inconsistent variable or in a constraint, and (3) all variables inside of q (not endpoints) are consistent in f. Note that if f is a

valid edge labeling for which there is an augmenting f-walk, then f is non-optimal (since $f \oplus q$ is a valid edge labeling with 1 or 2 fewer inconsistent variables).

The main result of this section is tractability of effectively coverable Δ-matroids.

Theorem (Theorem 9 restated)**.** *Given an edge CSP instance I with effectively coverable Δ-matroid constraints, an optimal edge labeling f of I can be found in time polynomial in $|I|$.*

The rough intuition of the algorithm for improving coverable Δ-matroid edge CSP instances is the following. When dealing with general Δ-matroids, augmenting f-walks may also end in a constraint – let us say that I has the augmenting f-walk q that ends in a constraint C. In that case, the parity of $f(D)$ and $(f \oplus p)(D)$ is the same for all $D \neq C$. If we guess the correct C (in fact, we will try all options) and flip its parity, we can, under reasonable conditions, find this augmentation via the algorithm for even Δ-matroids.

Not all Δ-matroids M are coverable. However, we will show below how to cover many previously considered classes of Δ-matroids. These would be co-independent [10], compact [15], local [7], and binary [13, 7] Δ-matroids.

Proposition 47. *The classes of co-independent, local, compact, and binary Δ-matroids are effectively coverable.*

The proof of this proposition as well as (some of) the definitions can be found at the end of this appendix.

C.1 The algorithm

The following lemma is a straightforward generalization of the result given in Lemma 30.

Lemma 48. *Let f, g be valid edge labelings of instance I (with general Δ-matroid constraints) such that g has fewer inconsistencies than f. Then we can, given f and g, compute in polynomial time an augmenting f-walk p (possibly ending in a constraint, in the sense of Definition 46).*

Proof. We proceed in two stages like in the proof of Lemma 30: First we modify g so that any variable consistent in f is consistent in g, then we look for the augmenting f-walk in $f \,\Delta\, g$. The only difference over Lemma 30 is that our g-walks and f-walks can now end in a constraint as well as in a variable.

First, we repeatedly modify the edge labeling g using the following procedure:

1. Pick a variable $v \in V$ which is consistent in f, but not in g. (If no such v exists then go to the next paragraph). By the choice of v, there exists a unique edge $\{v, C\} \in f \,\Delta\, g$. If $g(C) \oplus v \in C$, replace g with $g \oplus vC$, then go to the beginning and repeat. Otherwise, pick variable $w \neq v$ in the scope of C such that $\{w, C\} \in f \,\Delta\, g$ and $g(C) \oplus v \oplus w \in C$ (it exists since C is a Δ-matroid and $g(C) \oplus v \notin C$). Replace g with $g \oplus (vCw)$ and then also go to the beginning and repeat.

It can be seen that g remains a valid edge labeling, and the number of inconsistencies in g never increases. Furthermore, each step decreases $|f \,\Delta\, g|$, so this procedure must terminate after at most $O(|\mathcal{E}|) = O(|V|)$ steps.

We now have valid edge labelings f, g such that f has more inconsistencies than g, and variables consistent in f are also consistent in g. In the second stage we will maintain an f-walk p and the corresponding valid edge labeling $f^\star = f \oplus p$. To initialize, pick a variable $r \in V$ which is consistent in g but not in f, and set $p = r$ and $f^\star = f$. We then repeatedly apply the following step:

2. Let v be the endpoint of p. The variable v is consistent in g but not in $f^\star$, so there must exist a unique edge $\{v, C\} \in f^\star \Delta g$. If $f^\star(C) \oplus v \in C$, then output pC (an augmenting f-walk). Otherwise, pick variable $w \neq v$ in the scope of C such that $\{w, C\} \in f^\star \Delta g$ and $f^\star(C) \oplus v \oplus w \in C$ (it exists since C is a Δ-matroid and $f^\star(C) \oplus v \notin C$). Append vCw to the end of p, and accordingly replace $f^\star$ with $f^\star \oplus (vCw)$ (which is valid by the choice of w). As a result of this update, edges $\{v, C\}$ and $\{w, C\}$ are removed from $f^\star \Delta g$.

   If w is inconsistent in f, then output p (which is an augmenting f-walk) and terminate. Otherwise w is consistent in f (and thus in g) but not in $f^\star$; in this case, go to the beginning and repeat.

It is easy to verify that the p being produced is an f-walk. Also, each step decreases $|f^\star \Delta g|$ by 2, so this procedure must terminate after at most $O(|\mathcal{E}|) = O(|V|)$ steps and just like in the case of even Δ-matroids, the only way to terminate is to find an augmentation. □

Definition 49. Let f be a valid edge labeling of instance I with coverable Δ-matroid constraints. For a constraint $C \in \mathcal{C}$ and a Δ-matroid $C' \subseteq C$, we will denote by $I(f, C, C')$ the instance obtained from I by replacing the constraint relation of C by C' and the constraint relation of each $D \in \mathcal{C} \setminus \{C\}$ by the even Δ-matroid $D_{f(D)}$ (that comes from the covering).

Observe that f induces a valid edge labeling for $I(f, C, C)$. Moreover, if we choose $\alpha \in C$, then $I(f, C, \{\alpha\})$ is an edge CSP instance with even Δ-matroid constraints and hence we can find its optimal edge labeling by Algorithm 1 in polynomial time.

Lemma 50. *Let f be a non-optimal valid edge labeling of instance I with coverable Δ-matroid constraints. Then there exist $C \in \mathcal{C}$ and $\alpha \in C$ such that the optimal edge labeling for $I(f, C, \{\alpha\})$ has fewer inconsistencies than f.*

Proof. If f is non-optimal for I, then by Lemma 48 there exists an augmenting f-walk q in I. Take q such that no proper prefix of q is augmenting (i.e. we can not end early in a constraint). Let C be the last constraint in the walk and let $\alpha = (f \oplus q)(C)$.

We claim that $f \oplus q$ is also a valid edge labeling for the instance $I(f, C, \{\alpha\})$. Since we choose α so that $(f \oplus q)(C) = \alpha$, we only need to consider constraints different from C. Assume that p is the shortest prefix of q such that $(f \oplus p)(D)$ is not reachable from $D_{f(D)}$ for some $D \neq C$ (if there is no such thing, then $(f \oplus q)(D) \in D_{f(D)}$ for all $D \neq C$). We let $p = p^\star x D y$. Since $(f \oplus p^\star)(D)$ is reachable from $f(D)$, but $(f \oplus p^\star)(D) \oplus x \oplus y$ is not, we must have $(f \oplus p^\star)(D) \oplus x \in D$. But then $p^\star x D$ is an augmenting f-walk in I that is shorter than p, a contradiction with the choice of p. □

Lemma 51. *Let f be a valid assignment for the instance I with coverable Δ-matroid constraints and let $C \in \mathcal{C}$ and $\alpha \in C$ be such that there exists a valid edge labeling g for the instance $I(f, C, \{\alpha\})$ with fewer inconsistencies than f. Then there exists an augmenting f-walk for I and it can be computed in polynomial time given g.*

Proof. We begin by noticing that both f and g are valid edge labelings for the instance $I(f, C, C)$ (still with g having fewer inconsistencies than f). Hence by Lemma 48 we can compute an f-walk q which is augmenting in $I(f, C, C)$. It is easy to examine q and check if some proper prefix of q is an augmenting f-walk for I (ending in a constraint). If that happens we are done, so let us assume that this is not the case. We will show that then q must be an augmenting f-walk for I.

First assume that every prefix of q with integral length is an f-walk in I. Then either q is of integral length and we are done (q is its own prefix), or q ends in a constraint. If it is the latter, q must end in C, since that is the only constraint of $I(f, C, C)$ that is not forced to be an even

Δ-matroid. But the constraint relation C is the same for both I and $I(f, C, C)$, so flipping the last edge of q is allowed in I.

Let now p be the shortest prefix of q with integral length which is not an f-walk in I. We can write $p = p^\star xDy$ for suitable x, y, D. The constraint relation of D must be different in I and $I(f, C, C)$, so $D \neq C$. By the choice of p, for any prefix r of $p^\star$ of integral length we have $(f \oplus r)(D) \in D$ and moreover the tuple $(f \oplus r)(D)$ is reachable from $f(D)$. (If not, take the shortest counterexample r. Obviously, $r = r^\star uDv$ for some variables u, v and a suitable $r^\star$. Since $(f \oplus r^\star)(D) \in D$ is reachable from $f(D)$ and $(f \oplus r^\star)(D) \oplus u \oplus v$ is not, we get $(f \oplus r^\star)(D) \oplus u \in D$ and $r^\star uD$ is augmenting in I, which is a contradiction.) This holds also for $r = p^\star$, so $(f \oplus p^\star)(D)$ is reachable from $f(D)$.

To finish the proof, let $\beta^\star = (f \oplus p^\star)(D)$ and $\beta = (f \oplus p)(D)$. We showed that $\beta^\star \in D$ is reachable from $f(D)$. Also, $\beta^\star \oplus x \oplus y = \beta \in D_{f(D)} \setminus D$. Then by the definition of coverable Δ-matroids we have $\beta^\star \oplus x \in D$. Thus $p^\star xD$ is an augmenting f-walk in I and we are done.

It is easy to see that all steps of the proof can be made algorithmic. □

Now the algorithm is very simple to describe. Set some valid edge labeling f and repeat the following procedure. For all pairs (C, α) with $\alpha \in C$ and $C \in \mathcal{C}$, call Algorithm 1 on the instance $I(f, C, \{\alpha\})$ (computing the instance $I(f, C, \{\alpha\})$ can be done in polynomial time because all constraints of I come from an effectively coverable class). If for some (C, α) we obtained an edge labeling of $I(f, C, \{\alpha\})$ with fewer inconsistencies, use Lemma 51 to get an augmenting f-walk for I. Otherwise, we have proved that the original f was optimal.

The algorithm is correct due to Lemma 50. The running time is polynomial because there are at most $|I|$ pairs (C, α) such that $\alpha \in C$ and at most $|I|$ inconsistencies in the initial edge labeling, so the (polynomial) Algorithm 1 gets called at most $|I|^2$ times.

C.2 Even-zebras are coverable (but not vice versa)

The paper [11] introduces several classes of zebra Δ-matroids. For simplicity, we will consider only one of them: $\mathcal{C}$-zebras.

Definition 52. Let $\mathcal{C}$ be a subclass of even Δ-matroids. A Δ-matroid M is a $\mathcal{C}$-zebra if for every $\alpha \in M$ there exists an even Δ-matroid M_α in $\mathcal{C}$ that contains all tuples in M of the same parity as α and such that for every $\beta \in M$ and every $u, v \in V$ such that $\beta \oplus u \oplus v \in M_\alpha \setminus M$ we have $\beta \oplus v, \beta \oplus u \in M$.

In [11], the authors show a result very much similar to Theorem 9, but for $\mathcal{C}$-zebras: In our language, the result states that if one can find optimal labelings for $\mathrm{CSP}_{\mathrm{EDGE}}(\mathcal{C})$ in polynomial time, then the same is true for the edge CSP with $\mathcal{C}$-zebra constraints. In the rest of this section, we show that coverable Δ-matroids properly contain the class of $\mathcal{C}$-zebras with $\mathcal{C}$ equal to all even Δ-matroids (this is the largest $\mathcal{C}$ allowed in the definition of $\mathcal{C}$-zebras) – we will call this class *even-zebras* for short.

Observation 53. *Let M be an even-zebra. Then M is coverable.*

Proof. Given $\alpha \in M$, we can easily verify that the matroids M_α satisfy all conditions of the definition of coverable Δ-matroids: Everything reachable from α has the same parity as α and the last condition from the definition of even-zebras is identical to coverability. □

Moreover, it turns out that the inclusion is proper: There exists a Δ-matroid that is coverable, but is not an even-zebra.

Let us take $M = \{(0,0,0), (1,1,0), (1,0,1), (0,1,1), (1,1,1)\}$ and consider $N = M \times M$. It is easy to verify that M is a Δ-matroid that is an even-zebra with the sets M_α equal to $\{(0,0,0), (1,1,0), (1,0,1), (0,1,1)\}$ and $\{(1,1,1)\}$, respectively, and thus M is coverable.

Since coverable Δ-matroids are closed under direct products, N is also coverable. However, N is not an even-zebra: Assume that there exists a set N_α that contains all tuples of N of odd parity and satisfies the zebra condition. Then the two tuples $(1,1,1,0,0,0)$ and $(1,1,0,1,1,1)$ of N belong to N_α. Since N_α is an even Δ-matroid, switching in the third coordinate yields that N contains the tuple $(1,1,0,1,0,0)$ (this is without loss of generality; the other possibilities are all symmetric). This tuple is not a member of N, yet we got it from $(1,1,1,0,0,0) \in N$ by switching the third and fourth coordinate. So in order for the zebra property to hold, we need $(1,1,1,1,0,0) \in N$, a contradiction.

The above example also shows that even-zebras, unlike coverable Δ-matroids, are not closed under direct products.

D Classes of Δ-matroids that are effectively coverable

As we promised, here we will show that all classes of Δ-matroids that were previously known to be tractable are effectively coverable.

D.1 Co-independent Δ-matroids

Definition 54. A Δ-matroid M is *co-independent* if whenever $\alpha \notin M$, then $\alpha \oplus u \in M$ for every u in the scope of M.

Let V be the set of variables of M. In this case we choose M_α to be the Δ-matroid that contains all members of $\{0,1\}^V$ of the same parity as α. This trivially satisfies the first two conditions in the definition of a Δ-matroid. To see the third condition, observe that whenever $\gamma \in M_\alpha \setminus M$, the co-independence of M gives us that $\gamma \oplus u \in M$ for every $u \in V$, so we are done.

Moreover, each set M_α is roughly as large as M itself: A straightforward double counting argument gives us that $M \geq 2^{|V|-1}$, so listing M_α can be done in time linear in $|M|$.

D.2 Compact Δ-matroids

We present the definition of compact Δ-matroids in an alternative form compared to [15].

Definition 55. Function $F\colon \{0,1\}^V \to \{0,\ldots,|V|\}$ is called a *generalized counting function* (gc-function) if

1. for each $\alpha \in \{0,1\}^V$ and $v \in V$ we have $F(\alpha \oplus v) = F(\alpha) \pm 1$ and;
2. if $F(\alpha) > F(\beta)$ for some $\alpha, \beta \in \{0,1\}^V$, then there exist $u, v \in \alpha \,\Delta\, \beta$ such that $F(\alpha \oplus u) = F(\alpha) - 1$ and $F(\beta \oplus v) = F(\beta) + 1$

An example of such function is the function which simply counts the number of ones in a tuple.

Definition 56. We say that a $S \subseteq \{0,1,\ldots n\}$ is 2-*gap free* if whenever $x \notin S$ and $\min S < x < \max S$, then $x+1, x-1 \in S$. A set of tuples M is *compact-like* if $\alpha \in M$ if and only if $F(\alpha) \in S$ for some gc-function F and a 2-gap free subset S of $\{0,1,\ldots|V|\}$.

The difference to the presentation in [15] is that they give an explicit set of possible gc-functions (without using the term gc-function). However, we decided for more brevity and omit the description of the set.

Lemma 57. *Each compact-like set of tuples M is a Δ-matroid.*

Proof. Let the gc-function F and the 2-gap free set S witness that M is compact-like. Take $\alpha, \beta \in M$ and $u \in \alpha \,\Delta\, \beta$. If $F(\alpha \oplus u) \in S$, then $\alpha \oplus u \in M$ and we are done. Thus we have $F(\alpha \oplus u) \neq F(\beta)$.

Let us assume $F(\alpha \oplus u) > F(\beta)$. Since F is a gc-function we can find $v \in (\alpha \oplus u) \,\Delta\, \beta$ (note that $u \neq v$) such that $F(\alpha \oplus u \oplus v) = F(\alpha \oplus u) - 1$. Now we have either $F(\alpha) = F(\alpha \oplus u \oplus v) \in S$, or $F(\alpha) > F(\alpha \oplus u) > F(\alpha \oplus u \oplus v) \geq F(\beta)$, which means $F(\alpha \oplus u \oplus v) \in S$ because S does not have 2-gaps.

The case when $F(\alpha \oplus u) < F(\beta)$ is handled analogously. □

It turns out that any practical class of compact-like Δ-matroids is effectively coverable:

Lemma 58. *Assume $\mathcal{M}$ is a class of compact-like Δ-matroids where the description of each $M \in \mathcal{M}$ includes a set S_M (given by a list of elements) and a function F_M witnessing that M is compact-like and there is a polynomial p such that the time to compute $F_M(\alpha)$ is at most $p(|M|)$. Then $\mathcal{M}$ is effectively coverable.*

Proof. Given $M \in \mathcal{M}$ and $\alpha \in M$, we let M_α be the compact-like even Δ-matroid given by the function F_M and the set $U = [\min S_M, \max S_M] \cap \{F_M(\alpha) + 2k \colon k \in \mathbb{Z}\}$. It is an easy observation that $\alpha, \beta \in \{0,1\}^V$ have the same parity if and only if $F_M(\alpha)$ and $F_M(\beta)$ have the same parity, so all members of M_α have the same parity. In particular M_α contains all $\beta \in M$ of the same parity as α. Moreover, the set U is 2-gap free, so M_α is an even Δ-matroid.

Let now $\gamma \in M_\alpha \setminus M$. Then $F_M(\gamma) \notin S_M$. Since S_M is 2-gap free and $F_M(\gamma)$ is not equal to $\min S_M$, nor $\max S_M$, it follows that both $F_M(\gamma)+1$ and $F_M(\gamma)-1$ lie in S_M. Therefore, $\gamma \oplus v \in M$ for any $v \in V$ by the first property of gc-functions.

It remains to show how to construct M_α in polynomial time. We begin by adding to M_α all tuples of M of the same parity as α. Then we go through all tuples $\beta \in M$ of parity different from α and for each such β we calculate $F_M(\beta \oplus v)$ for all $v \in V$. If $\min S < F(\beta \oplus v) < \max S$, we add $\beta \oplus v$ to M_α. By the argument in the previous paragraph, this procedure will eventually find and add to M_α all tuples γ such that $F_M(\gamma) \in U \setminus S_M$. □

D.3 Local and binary Δ-matroids

We will avoid giving the definitions of local and binary Δ-matroids. Instead, we will rely on a result from [7] saying that both of these classes avoid a certain substructure. This will be enough to show that both binary and local matroids are effectively coverable.

Definition 59. Let M, N be two Δ-matroids where $M \subseteq \{0,1\}^V$. We say that M contains N as a *minor* if we can get N from M by a sequence of the following operations: Choose $c \in \{0,1\}$ and $v \in V$ and take the Δ-matroid we obtain by fixing the value at v to c and deleting v:

$$M_{v=c} = \{\beta \in \{0,1\}^{V \setminus \{v\}} \colon \exists \alpha \in M,\, \alpha(v) = c \wedge \forall u \neq v,\, \alpha(u) = \beta(u)\}.$$

Definition 60. The interference Δ matroid is the ternary Δ-matroid given by the tuples $\{(0,0,0)$, $(1,1,0)$, $(1,0,1)$, $(0,1,1)$, $(1,1,1)\}$. We say that a Δ-matroid M is *interference free* if it does not contain any minor isomorphic (via renaming variables or flipping the values 0 and 1 of some variables) to the interference Δ-matroid.

Lemma 61. *If M is an interference-free Δ-matroid and $\alpha, \beta \in M$ are such that $|\alpha \,\Delta\, \beta|$ is odd, then we can find $v \in \alpha \,\Delta\, \beta$ so that $\alpha \oplus v \in M$.*

Proof. Let us take $\beta' \in M$ so that $\alpha \Delta \beta' \subseteq \alpha \Delta \beta$ and $|\alpha \Delta \beta'|$ is odd and minimal possible. If $|\alpha \Delta \beta'| = 1$, we are done. Assume thus that $|\alpha \Delta \beta'| = 2k + 3$ for some $k \in \mathbb{N}_0$. Applying the Δ-matroid property on α and β' (with α being the tuple changed) k many times, we get a set of $2k$ variables $U \subseteq \alpha \Delta \beta'$ such that $\alpha \oplus U \in M$ (since β' is at minimal odd distance from α, in each step we need to switch exactly two variables of α).

Let the three variables in $\alpha \Delta \beta' \setminus U$ be x, y, and z and consider the matroid P on x, y, z we get from M by fixing the values of all $v \notin \{x, y, z\}$ to those of $\alpha \oplus U$ and deleting these variables afterward. Moreover, we switch 0s and 1s so that the triple corresponding to $(\alpha(x), \alpha(y), \alpha(z))$ is $(0, 0, 0)$. We claim that P is the interference matroid: It contains the triple $(0, 0, 0)$ (because of $\alpha \oplus U$) and $(1, 1, 1)$ (as witnessed by β') and does not contain any of the triples $(1, 0, 0)$, $(0, 1, 0)$, or $(0, 0, 1)$ (for then β' would not be at minimal odd distance from α). Applying the Δ-matroid property on $(1, 1, 1)$ and $(0, 0, 0)$ in each of the three variables then necessarily gives us the tuples $(0, 1, 1)$, $(1, 0, 1)$, and $(1, 1, 0) \in P$. □

Corollary 62. *Let M be an interference-free Δ-matroid. If M contains at least one even tuple then the set* $\mathrm{Even}(M)$ *of all even tuples of M forms a Δ-matroid. The same holds for* $\mathrm{Odd}(M)$ *the set of all odd tuples of M. In particular, M is effectively coverable by the even Δ-matroids* $\mathrm{Even}(M)$ *and* $\mathrm{Odd}(M)$.

Proof. We show only that $\mathrm{Even}(M)$ is a Δ-matroid; the case of $\mathrm{Odd}(M)$ is analogous and the covering result immediately follows.

Take $\alpha, \beta \in \mathrm{Even}(M)$ and let v be a variable v such that $\alpha(v) \neq \beta(v)$. We want $u \neq v$ so that $\alpha(u) \neq \beta(u)$ and $\alpha \oplus u \oplus v \in M$. Apply the Δ-matroid property of M to α and β, changing the tuple α. If we get $\alpha \oplus v \oplus u \in M$ for some u, we are done, so let us assume that we get $\alpha \oplus v \in M$ instead. But then we recover as follows: The tuples $\alpha \oplus v$ and β have different parity, so by Lemma 61 there exists a variable u so that $(\alpha \oplus v)(u) \neq \beta(u)$ (i.e. $u \in \alpha \Delta \beta \setminus \{v\}$) and $\alpha \oplus v \oplus u \in M$. □

It is mentioned in [7] (Section 4) that the interference Δ-matroid is among the forbidden minors for both local and binary (minors B1 and L2) Δ-matroids. Thus both of those classes are effectively coverable.

References

[1] Libor Barto, Marcin Kozik, and Todd Niven. The CSP dichotomy holds for digraphs with no sources and no sinks (A positive answer to a conjecture of Bang-Jensen and Hell). *SIAM Journal on Computing*, 38(5):1782–1802, 2009.

[2] André Bouchet. Multimatroids I. coverings by independent sets. *SIAM Journal on Discrete Mathematics*, 10(4):626–646, 1997.

[3] Andrei Bulatov. A dichotomy theorem for constraints on a 3-element set. In *FOCS'02*, pages 649–658, 2002.

[4] Andrei Bulatov. Complexity of conservative constraint satisfaction problems. *ACM Trans. Comput. Logic*, 12(4):24:1–24:66, July 2011.

[5] Andrei Bulatov, Peter Jeavons, and Andrei Krokhin. Classifying the complexity of constraints using finite algebras. *SIAM Journal on Computing*, 34(3):720–742, 2005.

[6] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Holant problems and counting CSP. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing*, STOC '09, pages 715–724, New York, NY, USA, 2009. ACM.

[7] Victor Dalmau and Daniel Ford. *Mathematical Foundations of Computer Science 2003: 28th International Symposium, MFCS 2003, Bratislava, Slovakia, August 25-29, 2003. Proceedings*, chapter Generalized Satisfiability with Limited Occurrences per Variable: A Study through Delta-Matroid Parity, pages 358–367. Springer Berlin Heidelberg, Berlin, Heidelberg, 2003.

[8] Zdeněk Dvořák and Martin Kupec. On planar Boolean CSP. In *ICALP '15*, volume 9134 of *Lecture Notes in Computer Science*, pages 432–443. Springer Berlin Heidelberg, 2015.

[9] Jack Edmonds. Path, trees, and flowers. *Canadian J. Math.*, 17:449–467, 1965.

[10] Tomás Feder. Fanout limitations on constraint systems. *Theoretical Computer Science*, 255(1–2):281–293, 2001.

[11] Tomás Feder and Daniel Ford. Classification of bipartite Boolean constraint satisfaction through delta-matroid intersection. *SIAM Journal on Discrete Mathematics*, 20(2):372–394, 2006.

[12] Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory. *SIAM Journal on Computing*, 28(1):57–104, 1999.

[13] James F. Geelen, Satoru Iwata, and Kazuo Murota. The linear delta-matroid parity problem. *Journal of Combinatorial Theory, Series B*, 88(2):377 – 398, 2003.

[14] John Hopcroft and Robert Tarjan. Efficient planarity testing. *J. ACM*, 21(4):549–568, October 1974.

[15] Gabriel Istrate. Looking for a version of Schaefer's dichotomy theorem when each variable occurs at most twice. Technical report, University of Rochester, Rochester, NY, USA, 1997.

[16] Bernard M. E. Moret. Planar NAE3SAT is in P. *SIGACT News*, 19(2):51–54, June 1988.

[17] Thomas J. Schaefer. The complexity of satisfiability problems. In *STOC'78*, pages 216–226, 1978.

[18] Edward P. K. Tsang. *Foundations of constraint satisfaction.* Computation in cognitive science. Academic Press, 1993.